Deep Space 1

Navigation: Primary Mission

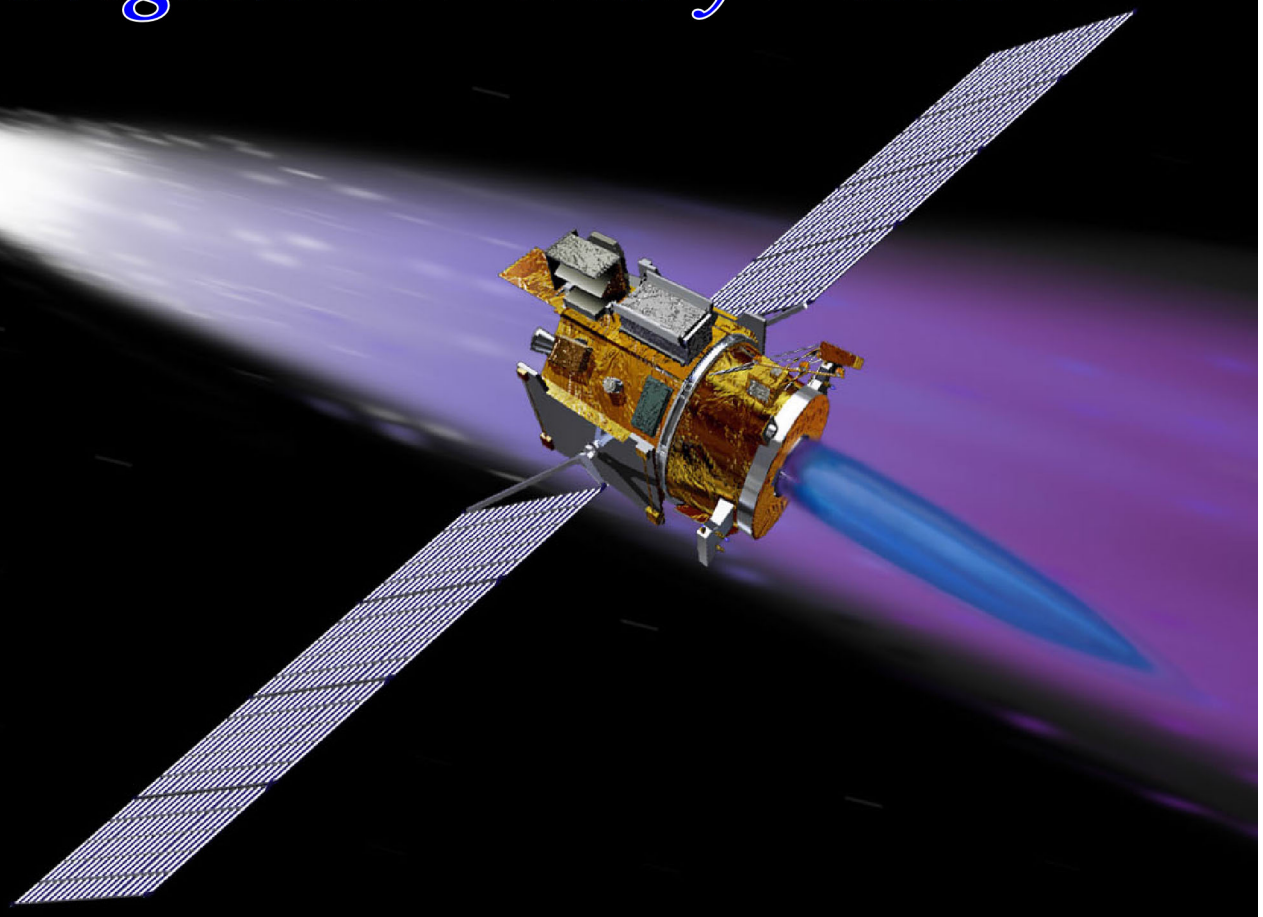

Brian Kennedy, J. Edmund Riedel, Shyam Bhaskaran, Shailen Desai, Don Han, Tim McElrath, George Null, Mark Ryne, Steve Synnott, Mike Wang, and Robert Werner

April 2004

JPL DESCANSO
Deep Space Communications and Navigation Systems Center of Excellence
Design and Performance Summary Series

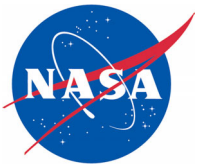

DESCANSO Design and Performance Summary Series

Article 8

Deep Space 1
Navigation: Primary Mission

Brian Kennedy, J. Edmund Riedel, Shyam Bhaskaran,
Shailen Desai, Don Han, Tim McElrath, George Null,
Mark Ryne, Steve Synnott, Mike Wang, and Robert Werner

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

National Aeronautics and
Space Administration
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

April 2004

This research was carried out at the
Jet Propulsion Laboratory, California Institute of Technology,
under a contract with the
National Aeronautics and Space Administration.

DESCANSO DESIGN AND PERFORMANCE SUMMARY SERIES

Issued by the Deep-Space Communications and Navigation Systems
Center of Excellence
Jet Propulsion Laboratory
California Institute of Technology

Joseph H. Yuen, Editor-in-Chief

Previously Published Articles in This Series

Article 1—"Mars Global Surveyor Telecommunications"
Jim Taylor, Kar-Ming Cheung, and Chao-Jen Wong

Article 2—"Deep Space 1 Telecommunications"
Jim Taylor, Michela Muñoz Fernández, Ana I. Bolea Alamañac, and Kar-Ming Cheung

Article 3—"Cassini Telecommunications"
Jim Taylor, Laura Sakamoto, and Chao-Jen Wong

Article 4—"Voyager Telecommunications"
Roger Ludwig and Jim Taylor

Article 5—"Galileo Telecommunications"
Jim Taylor, Kar-Ming Cheung, and Dongae Seo

Article 6—"Odyssey Telecommunications"
Andre Makovsky, Andrea Barbieri, and Ramona Tung

Article 7—"Deep Space 1 Navigation: Extended Missions"
Brian Kennedy, Shyam Bhaskaran, J. Edmund Riedel, and Mike Wang

Table of Contents

List of Figures

List of Tables

Foreword

This Design and Performance Summary Series, issued by the Deep Space Communications and Navigation Systems Center of Excellence (DESCANSO), is a companion series to the DESCANSO Monograph Series. Authored by experienced scientists and engineers who participated in and contributed to deep-space missions, each article in this series summarizes the design and performance for major systems such as communications and navigation, for each mission. In addition, the series illustrates the progression of system design from mission to mission. Lastly, it collectively provides readers with a broad overview of the mission systems described.

Joseph H. Yuen
DESCANSO Leader

Preface

This article will describe the means by which JPL navigators determined and controlled the trajectory of the Deep Space 1 (DS1) spacecraft during the course of its primary mission.

It is worth noting that DS1 was a fairly unconventional spacecraft. As such, the scope of this article will encompass unconventional methods of navigation (using autonomous acquired optical observations) as well as conventional methods (using radiometrics). This being the case, this article will discuss heavily the development, operation, and validation of the onboard optical navigation system. The interaction of this new navigation system with the navigation and science camera will be discussed, as will its interaction with the low-thrust ion propulsion system. The role played by conventional navigation in assessing the performance of both the new navigation system and the ion propulsion system will also be described.

It is hoped that this article can be used as a source which documents the navigation challenges (and rewards) of DS1, and that future navigators can gain insight into the implemention of new methods of guiding spacecraft towards scientifically rich encounters.

Acknowledgements

Special thanks are due to a number of individuals without whom AutoNav could not have been developed and proven. Marc Rayman, Phil Varghese, Leslie Livesay, and Dave Lehman, as leaders of the DS1 project, were continuously supportive and technically very contributing to the success. Steve Williams and Pam Chadbourne of the Mission Design Team made critical inputs to the navigation effort from the standpoint of mission design and planning. The DS1 software team led by Dan Eldred helped the AutoNav team as it ventured for the first time into the realm of deep-space flight software. The Testbed Team, led by Paula Pingree and Greg Harrison, made extensive testing of the AutoNav systems possible. Outstanding support in the sequencing area was provided by Curt Eggemeyer, Kathy Moyd, and "P.J." Guske. In the Navigation area, many thanks are due to Don Yeomans, Alan Chamberlin, and Mike Keesey for asteroid ephemeris development. Bill Owen was of tremendous help in taking ephemeris observations (especially of Braille), helping with the calibration of MICAS, and providing star catalogs. Finally, and perhaps most importantly, thanks go to the ACS team, led by Sima Lisman, and including Dan Chang, Tony Vanelli, Steve Collins, Sam Sirlin, Guru Singh, and John Esmiller.

Section 1

Mission and Spacecraft Description

1.1 Mission Overview

1.1.1 Technology Validation

On October 24, 1998, Deep Space 1 (DS1) became the first spacecraft launched under the New Millennium Program (NMP) [1]. The purpose of the NMP was to develop and certify new high-risk technologies for use in future low-cost science missions. DS1 served as an in-flight testbed for 12 new technologies. Beginning with those that involved or impacted navigation (Nav), the new technologies are as follows:

- Autonomous Onboard Optical Navigation (AutoNav)
- Miniature Integrated Camera and Spectrometer (MICAS)
- Solar Electric Ion Propulsion System (IPS)
- Solar Concentrator Arrays
- Miniature Integrated Ion and Electron Spectrometer
- Small Deep-Space Transponder
- Ka-Band Solid-State Power Amplifier
- Beacon Monitor Operations Experiment
- Remote Agent Experiment
- Low-Power Electronics
- Power Actuation and Switching Module
- Multi-Functional Structure

The concept of developing and validating new technologies in the context of a low-cost deep-space planetary mission was an extremely challenging one. In practice, the challenges were even greater.

Over the 11 months that comprised the DS1 primary mission, the DS1 flight team carried out intensive operations to characterize and validate all of these technologies, including AutoNav. The complete manifest of technologies was validated, with most of them proving highly successful.

Of these 12 technologies, three are of specific interest to this article: AutoNav, MICAS, and the IPS. The success of the onboard optical navigator hinged on the performance of the MICAS camera, which was AutoNav's sole source of observations. Also, a functional and well-behaved IPS was needed if AutoNav was to calculate and command successful changes to the mission trajectory during cruise phase and early in the encounter phase of the mission.

1.1.2 Mission Trajectory

Figure 1-1 shows the overall primary mission trajectory of DS1. While the primary goal of DS1 was the successful test and validation of the 12 technologies, a flyby of asteroid 19992KD (Braille) was scheduled in July 1999 as a "science bonus." To get to Braille at the

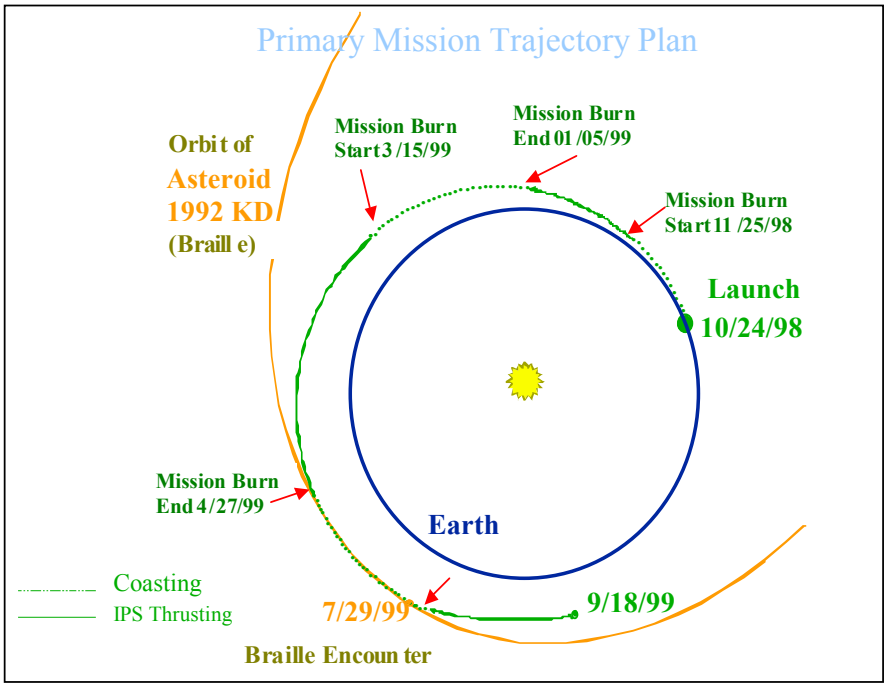

Fig. 1-1. Primary mission trajectory plan.

desired time and with the correct trajectory to make scheduled follow-on encounters during a potential extended mission, DS1 had to spend two and a half months thrusting with the IPS during the nine months between launch and encounter. Also, achieving this flyby required that the scheduled two and a half months of IPS thrusting begin within one month of launch. Beginning an extended IPS operation (referred to as "a mission burn") this early in the mission required that the onboard systems be fully validated within the first few weeks after launch. This underscores the intense nature of mission operations that defined the DS1 experience.

1.2 Spacecraft Overview

1.2.1. Spacecraft Configuration

The DS1 spacecraft (see Figure 1-2) consists of a roughly cylindrical bus with two large, gimbaled solar panels extending from each side (along the spacecraft Y-axis). The IPS thruster is mounted on the –Z-end of the spacecraft bus, with the camera and star tracker fields of view on the opposite side of the IPS (in the +Z-direction). In addition to the IPS, the DS1 spacecraft includes a monopropellant hydrazine Reaction Control System (RCS) for attitude control, consisting of eight 1-Newton thrusters. All of these are mounted close to the IPS, with four thrusting in the spacecraft +Z-direction (the same direction as the IPS thrust), and two each in the spacecraft +X and –X-directions. Since radiometric navigation depends on accurate spacecraft non-gravitational force modeling, it is important to note that the Z-facing thrusters (used for X- and sometimes Y-axis attitude control) operate in a completely unbalanced mode, while the X-facing thrusters have a balanced mode for Z-axis control, and a semi-balanced mode for Y-axis control. The X-thrusters are always used for turns about the Y-axis, since they have a much larger moment arm, but the Z-thrusters are used when finer Y-axis control is needed. Since there are no momentum wheels on the spacecraft, nominal attitude control is conventional "deadband" or "limit cycling" in nature, although the IPS (when running) provides attitude

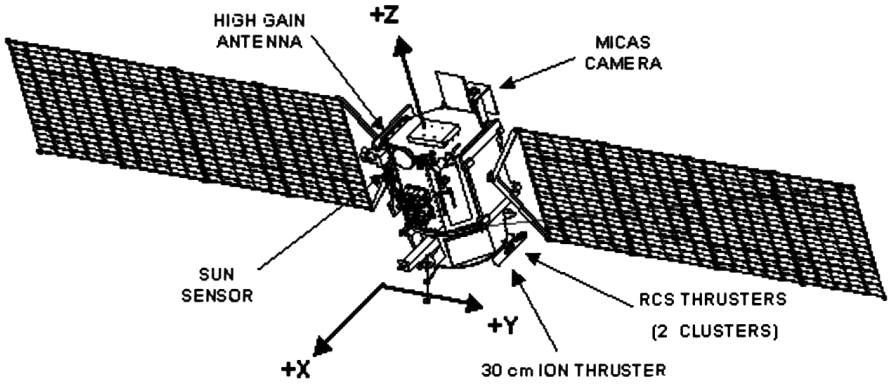

Fig. 1-2. The DS1 spacecraft.

control in the spacecraft X- and Y-axes through continuous gimballing. DS1 has three low-gain antennas (LGAs) with boresights in the spacecraft +Z, –Z, and +X-directions, and a high-gain antenna (HGA) also boresighted in the +X-direction.

1.2.2 Spacecraft Activities

During its mission, the DS1 spacecraft went through many activities that had a direct impact on navigation. As mentioned in Section 1.2.1, the spacecraft camera, HGA, and IPS thruster were mounted along fixed spacecraft axes (+Z, +X, and –Z, respectively). The need to align their respective boresights along various inertial attitudes, combined with a lack of reaction wheels, required that the Attitude Control Subsystem (ACS) use the unbalanced RCS thrusters to change the spacecraft attitude for each activity. During nominal operations, attitude changes would occur several times per week: once for HGA alignment during a high-rate data pass, once for thruster alignment to commence a week-long IPS thrust, and several times during a short (few-hour) period in which optical beacons were imaged for use as optical observations by the AutoNav system. Modeling of these activities was handled differently by the optical and radio navigators and will be discussed in Sections 3.1.1, 3.2.1.3, and 3.2.2.2.

1.3 Navigation Overview

Navigation for this mission was performed using conventional radiometric navigation techniques, as well as the onboard AutoNav system. It was acknowledged that the higher functions of AutoNav (picture taking and processing, orbit determination, etc.) could not be immediately relied upon, as they were parts of a new technology that had to be validated. Furthermore, even if fully operable and immediately invoked, AutoNav was not capable of performing the higher-accuracy near-Earth navigation (from immediately after launch to launch plus two days) required to assess injection conditions and keep the highly spacecraft-position-dependent near-Earth Deep Space Network (DSN) tracking within specification. Consequently, "conventional" radio navigation guided DS1 "out of the harbor" and, in fact, continued for the entire 1992KD (asteroid Braille) cruise, being the only independent means of assessing AutoNav

orbit determination (OD) performance. (As the actual mission proceeded, there was much dependence upon the Radio Navigation function as AutoNav was validated, but more importantly, as various problems with other subsystems were resolved or workarounds attempted.) Developing radio navigation techniques for use with a low-thrust mission was a technology development in and of itself.

1.3.1 Radio Navigation

The layout of DS1 and the attitude control actuators posed challenges to its navigation. The lack of either balanced thrusters or momentum wheels in the DS1 attitude control system configuration contributed significant challenges to all phases of radio navigation.

1.3.1.1 Radio Data Types. As with all National Aeronautics and Space Administration (NASA) deep-space missions, radiometric data are acquired during tracking passes, using antennas at the DSN complexes at Goldstone, Canberra, and Madrid. For DS1, conventional Doppler and range data acquired during these tracking passes were the sole data types used for orbit determination. It is worth pointing out that delta differential one-way range (DDOR) data acquisition was not used during any phase of the primary mission (but was used subsequently in the extended mission on approach to Borrelly).

1.3.1.1.1 Tracking Through the HGA. During a high-rate DSN pass, the ground communicated with the spacecraft through the spacecraft HGA while the spacecraft was at an Earth-pointing attitude. Since the HGA boresight was fixed along a spacecraft axis (specifically, the +X-axis), no IPS thrusting activity could be scheduled during the track without off-pointing the thrust vector from the desired thrust direction, which was, in general, undesirable.

1.3.1.1.2 Tracking Through One of the LGAs. During a low-rate tracking pass, the ground communicated with the spacecraft through one of the LGAs, while the spacecraft was at an IPS burn attitude. Due to the occasional use of smaller DSN antennas and the fairly weak LGA, telemetry was rarely available, even at low bit rates. During these passes only a limited amount of Doppler (2–3 hours) was received, but it provided a very strong visibility into the burn activity. With the absence of telemetry during these tracks, the provided Doppler signal quickly assessed evidence of the health of the spacecraft and its trajectory. Ranging data were typically not available during these tracks.

1.3.1.2 Modeling of Non-Gravitational Forces. The primary spacecraft non-gravitational perturbation models needed to navigate DS1 were solar radiation pressure (SRP), IPS thrusting, and RCS activity caused by turns and deadbanding. The modeling of SRP was fairly straightforward, since the largest force was imparted to the spacecraft through the solar panels. Due to the design of the solar panels and the yoke gimbal, the panels remained effectively perpendicular to the Sun at all times. Modeling of RCS and IPS activity is described in Section 3.1.1.

1.3.2 Optical Navigation Overview

The DS1 navigation system was the first use of autonomous navigation in deep space. The task for this system was to 1) perform interplanetary cruise orbit determination using images of distant asteroids, 2) control and maintain the orbit of the spacecraft using the ion propulsion system (another technology never before applied to deep space) and conventional thrusters,

3) perform approach orbit determination and control using images of the science targets, and 4) perform late knowledge updates of target position during close, fast flybys to facilitate a high degree of quality data return from 2 targets: asteroid Braille and comet Borrelly. Several functional components are necessary to accomplish these tasks. These include picture planning and image processing, dynamical modeling and integration, planetary ephemeris and star catalog handling, orbit determination data filtering and estimation, maneuver estimation, spacecraft ephemeris updates and maintenance, and general interaction with the other onboard autonomous systems.

1.3.2.1 Programmatic Components. The AutoNav system, as developed, was a complex piece of software. It was primarily a file-based computational system acting in response to sequenced commands from the ground operators. The file components, Nav process descriptions, and interactions with other aspects of the spacecraft hardware and flight software are discussed fully in Section 4.

1.3.2.2 Validation. Aside from comparisons with radio navigation solutions, validation of the AutoNav technology required careful calibrations of the IPS and the MICAS Camera—two subsystems that were closely interfaced with the AutoNav system. The first function of AutoNav, that of orbit determination, depended entirely on processing of optical images taken using the MICAS camera. The second function of AutoNav, that of trajectory control, required several interfaces to the IPS. Effective pointing of the camera and thruster subsystems also required interfacing with the attitude control subsystem. These interfaces are discussed in Section 4.2. Calibration of the MICAS subsystem is discussed in Section 4.2.1.

Section 2

General Consideration of DS1 and Low-Thrust Navigation

2.1 Radio

DS1 radio navigation was performed using conventional Doppler and range observations. However, the low-thrust nature of DS1's IPS created some difficulty in performing orbit determination well enough to validate the performance of both the AutoNav system and the IPS. Studies conducted before and after launch (based on performance of the spacecraft) indicated there would be difficulties providing the level of OD performance needed to validate the onboard systems.

Radio tracking data consisting of Doppler measurements from a single station were shown by [3] to provide the geocentric range rate, right ascension, and declination of the spacecraft in a single tracking pass, due to the station's motion caused by the rotation of the Earth. Analytical estimates of the accuracy of the derived angular measurements originally neglected the effect of geocentric acceleration uncertainty, since the first several generations of spacecraft had fairly low non-gravitational errors compared to the tracking data accuracy.[1] This analysis was extended by [2] to include acceleration uncertainty for Doppler data, in addition to looking for any changes necessitated by the then unique high declinations experienced in tracking the Ulysses spacecraft in 1994–1995. The conclusion of [2] was that high levels of geocentric acceleration uncertainty result in a significant loss of declination information.

This finding is relevant to DS1 both as a low-thrust mission and as a spacecraft with a high level of unbalanced attitude-control acceleration. The maximum thrust of the IPS (coincident with the maximum power available on the mission trajectory) was about 78 milli-Newtons (mN), producing an acceleration of 1.6×10^{-7} km/s^2 for the 486-kg launch mass of DS1. The pre-flight IPS thrust magnitude was uncertain by up to 2 percent, but even 0.1 percent is still larger than the 10^{-10} km/s^2 acceleration uncertainties considered in [2] for Ulysses. Even without IPS operation, the DS1 Z-facing thrusters produce an average acceleration of up to 4×10^{-10} km/s^2, with short-term variations at 10 percent of that level. While balanced thrusters, reaction wheels, spinning spacecraft, or inertial estimates of spacecraft delta-V can reduce this effect for coasting spacecraft, any low-thrust spacecraft is very likely to have navigationally significant acceleration errors. Being able to perform the necessary validation of the AutoNav system and perform general OD, as well as assess the performance of the IPS, required estimation of these forces to a fairly high degree of precision.

These high uncertainties also affected the information collected using range data. To illustrate the effect on orbit accuracy, Figure 2-1 shows angular uncertainties as a function of pass length and acceleration for a declination of 10 deg. The data used for this figure are based on 0.01 mm/s Doppler data every minute and 1.0-m range every five minutes. Due to the symmetry in the problem, acceleration uncertainty does not affect the right ascension uncertainty as declination uncertainty (as shown in Fig. 2-1). Therefore, it is concluded that right ascension

[1] A notable exception to this was the Voyager 1 spacecraft, which did have large non-gravitional impulses from its RCS. The use of near-simultaneous ranging and stochastic non-grav estimation during/around a 0-deg declination crossing in 1980 is described in [9].

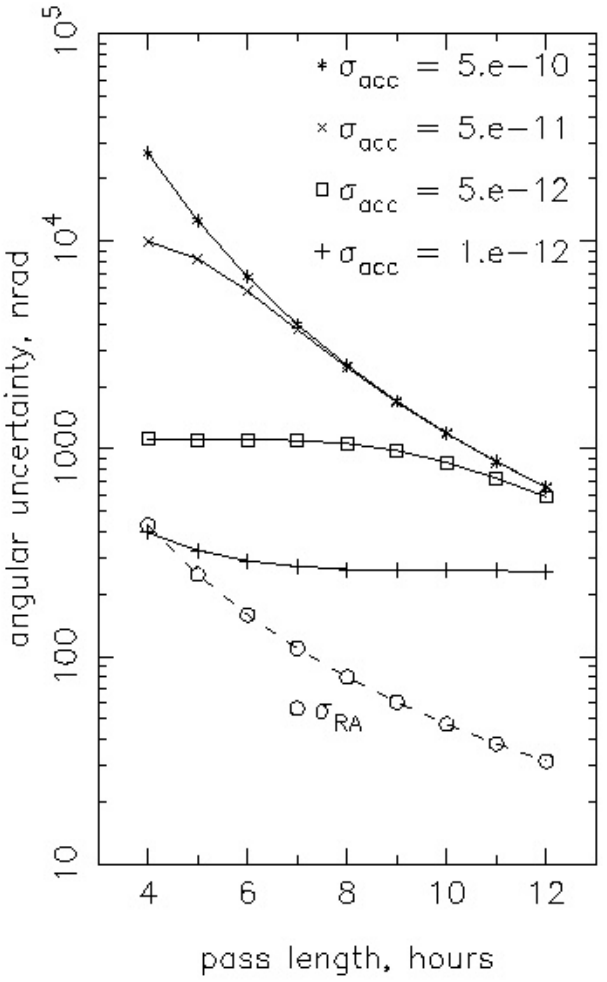

Fig. 2-1. Geocentric angular uncertainties as a function of acceleration uncertainty and pass length. Declination values are connected by solid lines, and right acsension values by a dashed line.

uncertainty is primarily a function of pass length. It was observed that a 1-m range weight is roughly equivalent to 1 mm/s Doppler for angular information content, but the 0.1 mm/s X-band Doppler that is routinely available from the NASA/Jet Propulsion Laboratory (JPL) Deep Space Network is much more effective than 1-m ranging. It is important to note that 4- to 6-hour passes produce a relatively ineffective measure of angular position in at least one component, and improvements of a factor of 2 or more are typical for 2-hour-pass-length extensions starting from 4- to 6-hour pass durations.

Large acceleration errors have a significant effect on the declination uncertainty for a single tracking pass. However, two consecutive passes should reduce the correlation between declination and acceleration such that a better angular estimate is obtained. Simple analytic results, based on taking the acceleration estimate from one pass and using it as *a priori* for the next, suggest that longer single passes are equivalent to two consecutive shorter passes of the same aggregate length as long as the two passes span separate DSN facilities. (But this is probably a pessimistic assumption since other correlations are not taken into account.)

A representative tracking scenario consisting of two consecutive 6.5-hour passes was evaluated numerically using the spacecraft state and IPS thrust as the only estimated parameters, for a thrust uncertainty of 1 mN (corresponding to an acceleration uncertainty of 2×10^{-9} km/s^2), both as a constant and as a stochastic parameter with a 1-hour update rate and a process noise scaled to match the constant uncertainty at 24 hours (i.e., $24^{1/2}$ mN). The following conclusions were made: 1) two passes are about 100 times better than 1 pass with a constant acceleration, 2) adding range data (versus Doppler-only) is about 10 times better for 2 passes, but does not provide much improvement for just one pass, and 3) stochastic accelerations at this level produce results that are 10 to 100 times worse than a single constant acceleration, and reduce the 2-pass benefit to merely a factor of 10.

2.2 Optical

Before the development of the DS1 AutoNav flight system began, a simulation system was created to test the concept of interplanetary cruise autonomous optical navigation [8]. The conceptual core of the AutoNav approach applied to interplanetary spacecraft is the use of asteroids, typically main-belt asteroids, as navigational beacons, to be observed against a field of stars with positions well determined astrometrically. Main-belt asteroids are the usual targets because of their large number and typically large size, when compared, for example, with near-Earth orbiting asteroids. For a typical inner Solar System mission trajectory, the range to the nearest visible asteroids is 1 to 1.5 AU. When the DS1 AutoNav was being conceived, requirements were levied on the camera system (Table 4-1) that required imaging capability (given nominal spacecraft stability) of unresolved objects of 12th magnitude. Such a sensitivity would enable the use of several dozen asteroid beacons at any time. Such a selection of targets would enable orbit determination to be accomplished anywhere inside the asteroid belt to a few hundred kilometers.

One important advantage of an all-optical-data orbit determination system was the insensitivity of the data type to high-frequency and small-velocity perturbations. With a velocity-measuring data type such as Doppler, this propulsion system posed substantial problems, as discussed in the previous section. These problems were dealt with by the radio navigation that was performed as part of the DS1 operations and validation task, but they did not have to be addressed by the onboard AutoNav system.

Fundamental to the OD subsystem is the modeled representation of the spacecraft flight path. This representation defines the nature and extent of the parameterization and accuracy possible in the system. The Navigator models the spacecraft motion with a numerical n-body integration, using major Solar System bodies as perturbing forces. Non-gravitational perturbations to the spacecraft trajectory included in the model include a simple spherical body solar-pressure model, a scalar parameter describing IPS engine thrust efficiency, and small accelerations in three spacecraft axes. A spherical-body solar-pressure model is sufficient because, for the majority of the time, the spacecraft will have its solar panels oriented toward the Sun. Even though the spacecraft can maintain this orientation with any orientation of the bus-body about the panel yoke axis, the panel orientation dominates the solar pressure effect by far.

During the cruise phase, the optical system was theoretically capable of taking 250-km measurements, assuming centerfinding accuracy of 0.1 pixels, which was determined theoretically and empirically from Galileo mission experience to be feasible, and depending on the available set of beacon asteroids. Over one week's time, that represents the capability of

measuring velocity to about 0.4 m/s, or accelerations to about 1.3 mm/s^2. The IPS engine is capable of delivering a maximum of about 0.1-Nt thrust, but on average will only be capable of half of that during the mission due to power restrictions. DS1 has a mass of about 420 kg, and therefore a typical inflight acceleration is about 120 mm/s^2. The IPS engine thrust is believed to be predictable to about one percent, or about 1.2 mm/s^2. It is clear then that long-frequency signatures in the IPS performance will be barely perceptible to the optical system in one week's time. These errors must be modeled. The capability of the Navigator IPS thrust noise model did not nearly meet the requirements of the ground radio navigation system, which had a 0.1-mm/s velocity sensitivity and a comparable acceleration sensitivity. Nevertheless, this level of modeling accuracy was perfectly adequate for the all-optical AutoNav system. However, coping with the noise in the engine performance was still the single most complicating factor in the flight OD algorithms for the optical system.

During periods of non-thrusting, and in the 20 days before a planned encounter, conventional trajectory correction maneuvers (TCMs) were scheduled. These assumed the use of the IPS, with the exception of the final two maneuvers, which were to be executed using the hydrazine thrusters of the ACS. Table 2-1 shows the TCM schedule, with expected and associated OD errors mapped to encounter at each TCM for the final 20 days of approach to the (then) planned target, McAuliffe. The algorithm used to compute these maneuvers is the same as that used for the IPS control algorithm. Necessarily, however, the maneuver solution is for only three parameters—the three components of delta-velocity. Another important difference between a reaction control system TCM and an IPS control is that the former occurs in a relatively short period of time, whereas IPS controls can take hours or days. In most cases, the applied maneuvers are expected to be small, on the order of 1 m/s or less, which for the IPS will take less than two hours.

Table 2-1: Approach TCM schedule with associated OD performance statistics.

Time to Encounter	Range to McAullife (km)	Downtrack Error (km)	Crosstrack Error (km)
-20 d	12.6E6	570	660
-10 d	6.3E6	138	27.3
-5 d	3.1E6	69	5.5
-2.5 d	1.6E6	54	2.5
-1.5 d	0.9E6	44	1.5
-1.0 d	630E3	42	1.2
-12 h	315E3	40.2	0.89
-6 h	157E3	40.1	0.55
-3 h	72E3	40.1	0.50

Section 3

Navigation Description

3.1 Radio

3.1.1 Modeling Methodology

To obtain accurate orbit determination estimates over time, correct modeling of RCS thruster activity (especially in the spacecraft Z-axis, which is also the IPS thrust direction) and the time variability of the IPS thrust was essential. While the AutoNav software accumulated an estimate of the RCS delta-V, known as the NonGrav history file, the resulting data were not very useful for radio navigation due to the 10 mm/s reporting threshold and 2-hour timeout rate. These settings were appropriately tuned for the onboard Autonav system. Although re-tuning was possible, the ground Radio Navigation team preferred to work directly with raw data from the Attitude Control System (ACS). Spacecraft telemetry containing attitude estimates, RCS pulse counts, and RCS thrust on-times were processed to obtain attitude-rate changes across each thruster firing event. Then, with knowledge of the spacecraft moments of inertia and thruster moment arms, the RCS delta-V's were calculated to within a few percent in most cases, based on the Doppler residuals resulting from the use of the delta-V model. This approach does not work well when too many pulses are fired within a time too short to obtain a separate attitude rate between each pulse. Consequently, after the attitude-derived delta-V's were included, additional impulsive delta-V's were added as necessary, based on Doppler residuals. Since the attitude rates and the antenna offset were big enough to produce noticeable Doppler residuals during deadbanding (but not IPS thrusting), the attitude-rate estimates were used to approximately model and remove this effect.

IPS thrusting was modeled as a finite-burn maneuver that allowed discrete thrust changes but did not allow the time of the thrust change to be estimated except implicitly by allowing the maneuver start time and duration to be estimated. Consequently, prefit adjustments of the time of each thrust change were performed, with the aid of a tool that estimated the time of a slope change in the Doppler residuals without the thrust change modeled. The thrust levels were also adjusted to obtain prefit Doppler residuals within a few mm/s so that the process noise assumptions in the estimation filter could be kept reasonably small. Note that the IPS thrust modeling and the RCS delta-V modeling process are iterative, as impulsive delta-V's only become apparent when the IPS thrust model is mostly complete. The IPS also has instances of thrust drop-outs when the high-voltage power supplies are autonomously turned off momentarily to clear high-voltage faults (events which are often confirmed by both Doppler residuals and telemetry). These were modeled as 0.2- to 0.3-mm/s impulsive delta-V's in the spacecraft –Z-direction, thus canceling part of the IPS +Z-thrust.

3.1.2 Estimation Methods

The estimation process used a batch-sequential epoch-state least-squares filter, with independent stochastic update times possible for each stochastic parameter. The standard estimation assumptions for all radiometric orbit solutions referenced in this article are given in Table 3-1, and the IPS assumptions are given in Table 3-2. The simplicity of the solar pressure

model is justified given the amount of other non-gravitational activity and the relatively even dimensions of the spacecraft bus, which is small compared to the solar array area. The Doppler and range weights are tighter than would generally be considered reasonable for a longer arc, but the high level of quickly changing model parameters makes weights close to the observed noise level acceptable. Constant spacecraft accelerations are used to account for remaining errors and biases in the RCS delta-V model that are not distinct enough to be modeled impulsively. The higher value of the stochastic acceleration *a priori* σ (or process noise) was used for three 3- to 5-minute time spans close to the IPS start on November 24. The short IPS thrust-magnitude update rate, varying from 2 to 10 minutes depending on data noise and IPS stability, was designed to absorb most of the short-period effects, since there was not sufficient information content in the Doppler data to estimate thrust direction at the same time. The IPS thrust-direction estimates were instead allowed to vary slowly and were checked for general validity, with resulting changes to other parameters as necessary to enforce this result. Finally, the RCS delta-V's were allowed to change in magnitude by 50 percent σ to absorb errors in the RCS delta-V modeling process. It should be noted that there was a much higher degree of time variability reflected in this model than that used in [4] in analyzing IPS calibration performance, which highlights the importance of obtaining actual flight data.

Table 3-1. Standard estimation assumptions.

Parameter	*A priori* σ
Earth Orientation:	
Pole	30 nrad
UT1	50 nrad
Troposphere:	
Dry	1 cm
Wet	4 cm
Ionosphere (X-Band):	
Night	1.1 cm
Day	5.6 cm
Solar Pressure	
Radial (Gr)	10%
Normal (Gx, Gy)	5% of radial
Station Locations (fully correlated):	
Spin radius	8 to 9 cm
Z-height	7 to 8 cm
Longitude	5 to 6 cm
Range bias per pass	5 m

Table 3-2. IPS1 estimation assumptions.

Estimated Parameters			
Parameter		*A priori* σ	
Impulsive maneuver components		0.5 mm/s	
Thrust start time		1 s	
Data Weights			
Data type		Weight (1σ)	
Doppler (60 s)		0.03 mm/s	
Range		0.3 m	
Estimated and Stochastic Parameters			
Parameter	Estimated *a priori* σ	Stochastic *a priori* σ	Update Rate/Correl. Time
Spacecraft Accelerations			
X-, Z-axis	5×10^{-12} km/s^2	10^{-11} km/s^2	1 hour/none
(short term)		10^{-9} km/s^2	1 min/none
Y-axis	10^{-12} km/s^2	5×10^{-12} km/s^2	6 hours/none
IPS Thrust			
Magnitude	5 mN	2 mN	2–10 min/none
Direction	2 degrees	1 deg	1 hour/6hours
RCS Δ-V scale	Not estimated	50%	per Δ-V

3.1.3 Thrusting During Tracking

The nominal DS1 thrust arc schedule includes a halt in thrusting once per week for a HGA communications pass—a requirement imposed by spacecraft geometry if the optimal thrusting direction is to be maintained. From these numerical and analytic results, it is clear that avoiding continuous IPS thrusting while tracking is beneficial to orbit determination, even with the unbalanced RCS accelerations. While continuous thrusting during infrequent communication passes would be possible at the lowest thrust level without serious effects on the overall mission efficiency, doing so would require significantly more or different radio tracking data to maintain reasonable orbit solution accuracy. Since DS1 is nominally tracked only twice a week during thrusting, the range data obtained from the weekly HGA pass are important in determining the centroid of the delta-V, which would not be well determined by Doppler data alone. The possibility of unexpected halts in IPS thrusting also drives the choice of the ranging data modulus parameter toward the largest values, since the geocentric range could be quite uncertain (easily thousands of kilometers) in such an eventuality.

3.2 Optical

The theoretical basis of AutoNav is a process in which images of asteroids (typically main-belt) are taken against the distant stars, and through the measured parallax, geometric information is inferred. This information is used in a dynamic filter to determine the spacecraft position and velocity, as well as parameters describing the performance of the IPS and solar pressure. With this information, corrections to the mission design, as described in the propulsion profile, are made and/or predictions for necessary trajectory correction maneuvers (TCMs) are computed.

The AutoNav system is a set of software elements that interact with the imaging, attitude control, and ion propulsion systems aboard DS1. The principal elements and functions of AutoNav are:

- NavRT—provides critical ephemeris information to other onboard subsystems, such as the Attitude Control System.
- NavExec—plans and executes various important Nav-related activities, such as image-taking and processing, Ion Propulsion System thrusting events, and TCMs.
- ImageProcessor—the image processing subsystem.
- OD—the orbit determination computation element.
- ManeuverPlanner—performs computations relative to the IPS events and the TCMs.

While these capabilities and underlying algorithms are described here, the technological details are covered in Section 4.

3.2.1 Core Capabilities

3.2.1.1 Ephemeris Services. This is the required function to provide various systems onboard (but chiefly ACS) information about the location of the spacecraft and any Solar System object of importance to the mission, such as Earth (for telecommunications purposes) and other Solar-System bodies for camera targeting.

3.2.1.2 Image Acquisition and Processing. AutoNav was able to coordinate and execute the tasks needed to acquire and process the images used for optical navigation (opnav). Given an *a priori* listing of desired optical targets and an allotted completion time, AutoNav planned a picture-taking session and coordinated turn sequencing with help from the Attitude Planning Expert (APE), a subsystem of the ACS.

As the spacecraft was turned from target to target and took pictures, AutoNav processed the images to obtain useful navigation data. There were three stages to this process. First, there was an initial coarse registration, wherein the *a priori* prediction of the location of objects in the field, good to 10–20 pixels, was refined to 1–2 pixels. Next, precision astrometry took place, where the locations of objects were determined to 0.1–0.25 pixels (see Section 3.2.2.1). Finally, by using only the star images as reference, the inertial attitude of the camera when the image was taken was computed. That information, plus the location of the target, was written to a file for later use by the OD process.

3.2.1.3 Orbit Determination. This is the purely computational function of reducing the suite of optical observations on the OD file to an estimated state of the spacecraft. Sub-elements of this function include numerical integration of the spacecraft position and velocity, as well as partial derivatives of the spacecraft state with respect to dynamic parameters. Estimation and filtering itself is a key function (see Section 3.2.2.2).

Following the computation of a new spacecraft orbit, AutoNav integrated a new spacecraft ephemeris, produced a Spacecraft Ephemeris file, and made this file available to Ephemeris Services (see Section 3.2.1.1).

3.2.1.4 Trajectory Control and Maneuver Planning. Following an OD event, AutoNav was typically asked to compute any needed course corrections using a mission burn or a TCM. Computational elements involved in this function included iterative trajectory integration to

compute *a priori* mistargeting and numerical partial derivatives for estimating correction parameters. These parameters were either elements of a discrete RCS or Ion Propulsion System TCM or the directional and duration parameters for an IPS mission burn (see Section 3.2.2.3). Additionally, the Maneuver Planner determined, through interaction with APE, whether a proposed TCM was "legal" in the context of spacecraft orientation constraints. If there was a violation, further interactions with APE decomposed the TCM into two allowed legs via a process called "vectorization."

Additional functions of AutoNav allowed for the execution of planned TCMs and mission burns using either the IPS or the RCS. Managing these tasks included evaluating navigation files to determine the proper direction and duration of the burning and the starting and termination of the burns. It also included evaluating ACS interaction to specify the spacecraft attitude and IPS interaction to manage various engine activities (e.g., starting, stopping, pressurizing, setting throttle levels, and safing the engine).

3.2.1.5 Encounter-Related Activities: Image Processing and Orbit Determination. Autonomous encounter activities are handled within AutoNav by the encounter operations subsystem, known as Reduced State Encounter Navigation (RSEN). When in effect, RSEN responds to any active pixel sensor (APS) images sent to AutoNav by the MICAS camera by reducing the incoming image data into optical observations. Given successfully generated results of the image-processing function described above, RSEN performs the reduced state orbit-determination operation and transmission of the data to ACS for target tracking. Figure 3-1 shows the expected accuracy of the RSEN system in downtrack (i.e., time of flight) on approach to Braille, given successful picture delivery and processing at each of the indicated data. Note that two different *a priori* errors were assumed, 10 and 20 seconds, representing 150 and 300 km of downtrack error, respectively. In fact, the actual error was probably closer to 300 km based on the ephemeris errors observed in the cross-track directions during the Braille approach. Figure 3-1 shows a continuous representation of the expected RSEN performance.

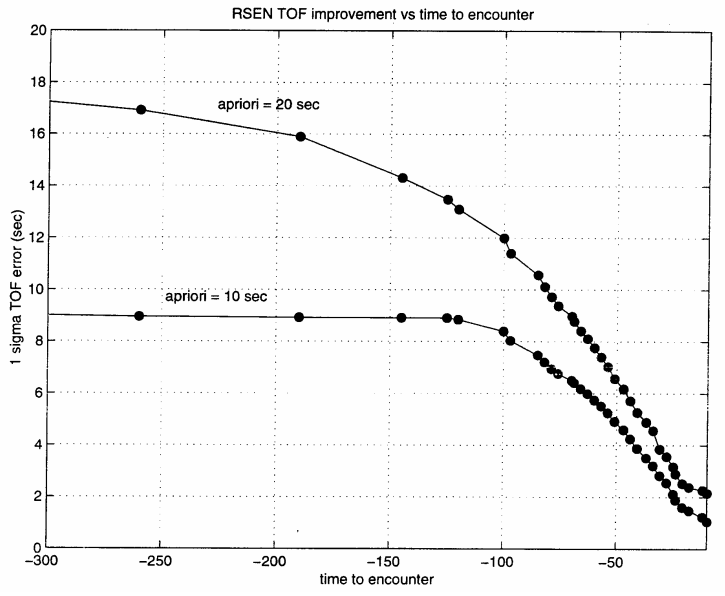

Fig. 3-1. RSEN time-of-flight performance.

The final step in the encounter process is to start encounter sequences at a time appropriate for encounter science data gathering. During a close flyby of the target, the acquisition of navigation knowledge about the relative downtrack position of the spacecraft happens only very late. Consequently, parts of the close approach science activity must be broken up into segments, generally getting shorter as they approach close-approach. Each of these segments is started independently with increasing time precision by RSEN. The function that starts the encounter sequences is completely dependent upon the computational processes outlined above for the determination of expected time of flight. Given this information, AutoNav, when asked to start an encounter sequence, immediately determines the time remaining to encounter and starts a mini-sequence to launch the desired sequence at the appropriate time.

3.2.2 Core Algorithm Descriptions

3.2.2.1 Multiple Cross-Correlation. Figure 3-2 shows a diagrammatic representation of the algorithm that forms the basis of the cruise image processing in AutoNav. The underlying assumption of the algorithm is that long exposures will be necessary to image dim objects, and thus, because of ambient motions of the spacecraft due to attitude maintenance by ACS, the images of stars and targets will be smeared, often in complicated patterns. These patterns, called "glyphs," will be nearly identical to one another, since the effects of "twisting" deadband motion in the field is small (the attitude maintenance is roughly equivalent in all directions, but maps to a much smaller effect in the field than the two cross line-of-sight pointing directions). Based on initial knowledge of pointing of the spacecraft (as provided by ACS) and predictions of the relative locations of the positions of objects in the field of view (based on the target ephemeris and the star catalog), segments of the pictures are extracted and normalized, and these become templates or "filters." Filters for each object are used to locate each of the other objects. Once all of the objects are located relative to one another (and these data filtered for bad or weak signal), a least-squares estimate is made of the relative offset of the objects relative to one another. A

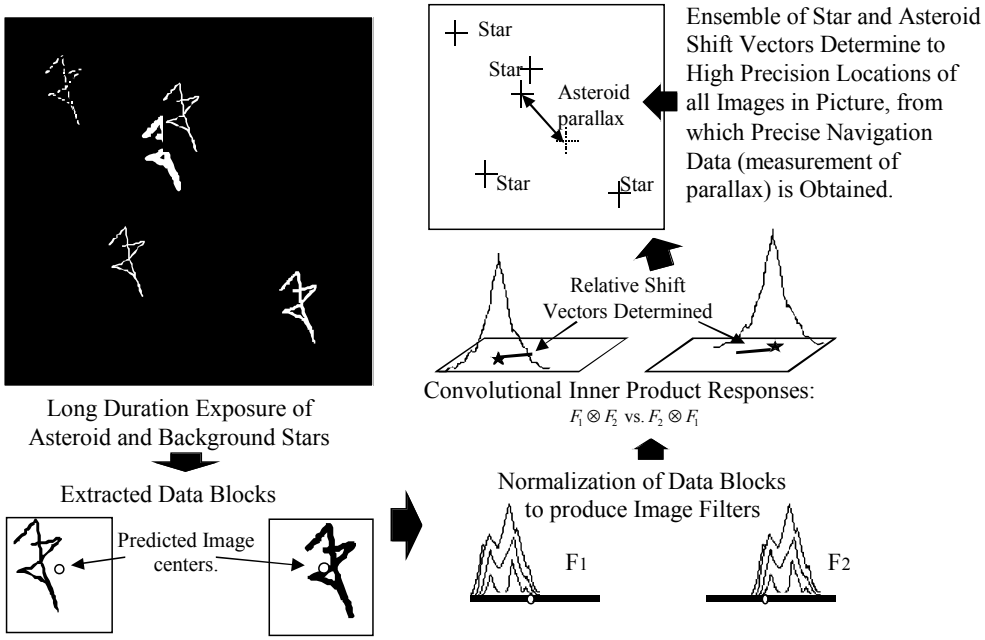

Fig. 3-2. Multiple cross-correlation of asteroid and stars.

complete description of this algorithm is given in [5], as it was used for the Galileo Gaspra encounter.

3.2.2.2 Orbit Determination. There are several crucial elements to the DS1 Orbit Determination function: 1) the numerical integration of the spacecraft trajectory, 2) the dynamic models of the gravitational and non-gravitational perturbations that drive that integration, 3) the generation and the mapping of the covariance in time, with the state transition matrix, and 4) the formation of the data filter itself. Reference [6] gives a complete development of the filter and related algorithms. The OD filter used is a Kalman batch-sequential least-squares filter. A typical data arc was planned to be about a month long, with four 1-week batches that corresponded to the typical one Photo-Op event per week. The estimated parameters for a given solution include the position and velocity at the beginning of the data arc, a constant-acceleration 3-vector which applies for the duration of the arc, and IPS thrust scale factors that are stochastic parameters for each week. The latter parameters were only estimated if there was an IPS mission burn in progress during that portion of the arc.

3.2.2.3 IPS Mission Burn Targeting. The process for retargeting the spacecraft trajectory during a mission burn is shown in Figure 3-3. This is an iterative application of a linear estimation of corrections to the direction of burn of an individual element of the multi-element mission burn and the duration of the final element. Since it is iterative, the overall algorithm is non-linear. The algorithm automatically decides how many segments to include in the solution, starting with a minimum acceptable number, and increasing the number as necessary to gain sufficient control authority to achieve convergence, i.e., putting the spacecraft on target.

X_e = Target encounter conditions ($B \cdot R, B \cdot T, tof$)

ΔX_e = Deviations from desired encounter

α_i = Rt. ascention of thrust segment i

δ_i = Declination of thrust segment i

τ_κ = Burn duration of final segment κ

$\Delta X_e = K \Delta s$, where...

$$K^T = \begin{bmatrix} (\partial X_e / \partial \alpha_{k_1}) \\ (\partial X_e / \partial \delta_{k_1}) \\ (\partial X_e / \partial \alpha_{k_1+1}) \\ (\partial X_e / \partial \delta_{k_1+1}) \\ \vdots \\ (\partial X_e / \partial \alpha_{k_2}) \\ (\partial X_e / \partial \delta_{k_2}) \\ (\partial X_e / \partial \tau_\kappa) \end{bmatrix}, \text{ and... } \Delta s = \begin{bmatrix} \Delta \alpha_{k_1} \\ \Delta \delta_{k_1} \\ \Delta \alpha_{k_1+1} \\ \Delta \delta_{k_1+1} \\ \vdots \\ \Delta \alpha_{k_2} \\ \Delta \delta_{k_2} \\ \Delta \tau_\kappa \end{bmatrix}$$

Then, by the calculus of variations,

$$\Delta s = K^T (K K^T)^{-1} \Delta X_e$$

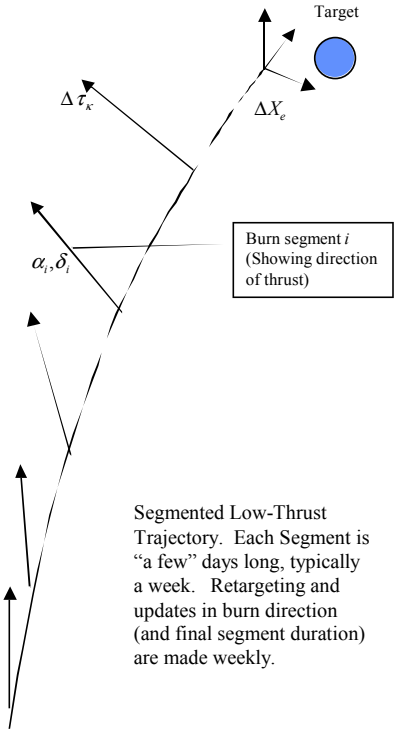

Fig. 3-3. Adjusting a low-thrust burn arc.

It is important to note that the spacecraft is given a converged trajectory initially. This trajectory has been discovered and reasonably converged initially with an algorithm known as “differential inclusion” [7] and uplinked to the spacecraft. Then, within well-regulated limits, the maneuver planner is allowed to adjust this trajectory to keep the spacecraft targeted.

Section 4

DS1 Optical Navigation Technology

4.1 AutoNav System Components

In order to perform many of the autonomous functions asked of it, AutoNav was a necessarily complex subsystem. The main computational components of the system are described here, including file elements, programmatic elements, and interactions with the ground operators and other subsystems onboard DS1.

4.1.1 AutoNav File Descriptions

AutoNav is a file-based computational system. Conditions necessary to operate AutoNav, such as operational parameters, planetary ephemerides, star catalog, etc., are provided by ground operators. This information provides AutoNav with sufficient information to start gathering its own data by scheduling and taking pictures. AutoNav updates these data as necessary as a means of storing computed information and/or communicating between the AutoNav links. Table 4-1 lists the AutoNav files and their update frequency (by AutoNav and the ground).

Table 4-1. AutoNav files, sizes, autonomy status, locations, and update frequency.

File Description	File Size (kB)	File Update Frequency		Location
		From Ground	Auto-Onboard	
Star Catalog	2200	1/mission	Never	EEPROM
Planetary Ephemeris	92	1/mission	Never	EEPROM
TCM Params	5	4/year	Never	EEPROM
Encounter (RSEN) Params	0.3	2/encounter	Never	EEPROM
Encounter Star Catalog	0.1	2/encounter	Never	EEPROM
FrankenKenny Params	0.7	2/encounter	Never	EEPROM
CCD Camera Params	0.6	2/year	Never	EEPROM
APS Camera Params	3	1/encounter	Never	EEPROM
Beacon Ephemeris File	2	2/year	Never	EEPROM
Mass Profile	56	4/year	Never	EEPROM
Picture plan	20	4/year	Never	EEPROM
Control Params	20	4/year	Never	EEPROM
Photo-Op Params	4	2/year	Never	EEPROM
IPSburn Params	0.4	2/year	Never	EEPROM
Nongrav Params	0.2	2/year	Never	EEPROM
Imageproc Params	0.3	2/year	Never	EEPROM
File of Filenames	1.5	4/year	1/month	EEPROM
Maneuver	33	4/year	Weekly	EEPROM
OD	10	2/year	Weekly	EEPROM
Spacecraft Ephemeris	12	1/year	Weekly	EEPROM
OpNav	1000	Never	Weekly	RAM
Nongrav History	40	Never	Several/day	EEPROM

4.1.1.1 Star Catalog. The Star Catalog (Starcat) file contains the positions and brightness of the stars necessary for navigation. For DS1, this file contained 220,000 stars in an annulus of ±30 deg of the ecliptic and as deep as 10.5 M. This catalog was extracted from a hybrid catalog composed of the Astrographic-Tycho Catalogue combined with Hipparcos data.

4.1.1.2 Planetary Ephemeris. This file contains the positions of nine planets and the moon represented as Chebyshev polynomials. This file extends for the duration of the primary and extended missions and is based on the JPL DE-403 planetary ephemeris.

4.1.1.3 TCM Params. The TCM Params file contains parameters that moderate the function of the TCM activities. These parameters include the minimum wait times between turns and actual burns of the RCS and IPS engines, and other such timing and control.

4.1.1.4 Encounter (RSEN) Params. This file contains parameters that regulate the activity of the close approach navigation system, called Reduced State Encounter Navigation.

4.1.1.5 Encounter Star Catalog. The Encounter Star Catalog file contains a small star catalog that is used only for the far encounter navigation image processing. A separate catalog is necessary to process the encounter pictures because of the geometry of the approach (e.g., outside the main catalog annulus) or because of the depth of stars necessary to include.

4.1.1.6 FrankenKenny Params. FrankenKenny (FK) is the onboard self-simulation subsystem of AutoNav. It creates images based (optionally) on an independent model of the spacecraft position and feeds these to AutoNav, providing closed-loop simulation. This file contains parameters to control the simulation.

4.1.1.7 CCD Camera Params. This file contains parametric descriptions of the MICAS charge-coupled device (CCD) camera, including focal length and distortion models.

4.1.1.8 APS Camera Params. This file is similar to the one above, but for the active pixel sensor (APS) visual channel of the MICAS camera.

4.1.1.9 Beacon Ephemeris. Beacon Ephemeris contains the Chebyshev polynomial description of several dozen asteroids used for navigation.

4.1.1.10 Mass Profile. The Mass Profile file contains a table of propellant consumption values—in essence, the predicted mass of the spacecraft at discrete times.

4.1.1.11 Picture Plan. This file contains recommended asteroid targets for particular Photo Opportunities, selected for maximum navigational strength and to minimize the amount of turn time required to move from target to target.

4.1.1.12 Control Params. Control Params contains dynamic modeling parameters for the spacecraft position integration and targeting parameters (such as the desired flyby conditions). This file also contains parameters used by the Orbit Determination routines.

4.1.1.13 Photo_Op Params. This file contains the parameters to control the Photo-Op operation, the navigation-controlled events that cause navigation images to be taken and processed. These parameters are primarily timing parameters such as delays after turns.

4.1.1.14 IPSburn Params. This file contains the parameters to control the operation of the Nav-directed "Mission Burns," which are long periods of IPS thrusting. These parameters are primarily timing parameters, e.g., delays after turns.

4.1.1.15 Non-grav Params. The Non-grav Params file contains parameters to direct the writing of the Nongrav History file, which is a continuous record of intentional non-gravitational events onboard accomplished by the ACS or IPS. These parameters largely regulate the precision in time with which this record is kept.

4.1.1.16 Imageproc Params. This file regulates the operation of the image processing operation, with controls such as thresholds for brightness and filtering gains.

4.1.1.17 File of Filenames. This is the navigation directory, containing the full path names of all of the navigation files, thereby indicating their locations in the file system. This file is automatically updated when files are updated using the Nav_Data_Update function.

4.1.1.18 Maneuver. This file contains the descriptions of thrusting events such as TCMs and mission burns. It also divides up time into segments for purposes of OD processing. The Maneuver file is autonomously updated by the AutoNav ManPlan function.

4.1.1.19 OD. The OD File contains the current best-estimated position of the spacecraft at several junctures in time through the data arc (typically a month). This file is autonomously updated during the Nav_Do_OD orbit determination function.

4.1.1.20 Spacecraft Ephemeris. Spacecraft Ephemeris is a Chebyshev Polynomial representation of the spacecraft position and velocity. This file is automatically updated after the Nav_Do_OD and Nav_ManPlan functions.

4.1.1.21 OpNav. This file contains the results of image processing in the Nav_Do_PhotoOp function: edited picture elements, and determined line/pixel positions of objects.

4.1.1.22 Nongrav History. This file contains the continuous record of intentional "non-gravitational" (nongrav) (i.e., thrusting) events onboard accomplished by the ACS or IPS.

4.1.2 Software System

The AutoNav software is shown in Figure 4-1. The AutoNav system is composed of three principal parts: Nav Executive, Nav Main, and Nav Real-Time (NavRT). These communicate with each other and other subsystems through the underlying system messaging facility. Much of the commanding by AutoNav is through the sequencing subsystem discussed below.

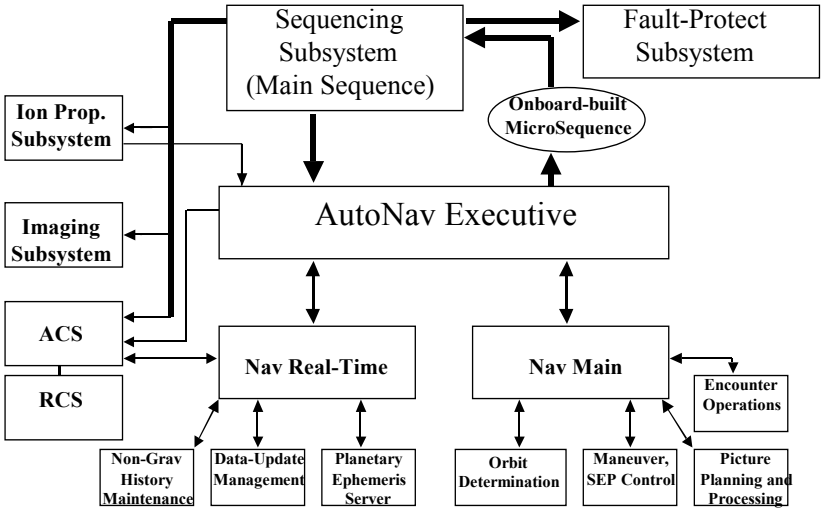

Fig. 4-1. The AutoNav software system and interacting system software.

4.1.2.1 Nav Executive. NavExec is AutoNav's director of spacecraft activities. It receives messages from other spacecraft subsystems and sends command directives, either through the onboard sequence machine or through direct messages to other subsystems. When using the sequence subsystem (sequence engine), NavExec will build small sequences and launch them. When NavExec needs an activity to occur immediately, for example, to turn the spacecraft to a desired burn attitude, it will build a relative time sequence that the sequence engine initiates at once. Alternatively, when NavExec needs to ensure that an event begins exactly at a certain time, it will build and initiate an absolute timed sequence, for example, to cause the main engine to ignite for a TCM. NavExec contains three main state machines for Photo-Ops, TCMs, and mission burns. These machines are mutually exclusive, since the activities involved are incompatible.

4.1.2.2 Nav Real-Time. NavRT is the subsystem of AutoNav that provides critical onboard ephemeris information to other onboard subsystems, but principally to ACS. NavRT operates at a much higher priority level in the flight software than the other AutoNav components due to the need to respond to sometimes frequent and time-critical ACS requests. NavRT also accomplishes file updates involving ephemeris related files by ensuring that changes in files are completed without jeopardizing ACS ephemeris queries.

4.1.2.3 Nav Main. Nav Main is the central computing element of AutoNav. Requests for activity that involve large amounts of computing are directed to Nav Main by NavExec or go to Nav Main directly through the command subsystem. These functions include picture processing requests from NavExec, Do-OD, and ManPlan commands from ground commands. There are several important subfunctions of Nav Main:

- Trajectory integration, which includes dynamic modeling of gravitational and non-gravitational forces acting on the spacecraft.
- Data filtering, including a U-D factorized batch-sequential filter.
- Trajectory update computation based on an iterative linear minimum-norm solution for changes to the IPS thrust profile to reduce projected targeting errors.

4.1.3 AutoNav Commanding Strategy

DS1 AutoNav is fully autonomous, but only upon invitation of ground controllers. Most importantly, AutoNav will cause physical spacecraft activity or intense computational action only when invited to do so by the ground, allowing controllers to be fully aware beforehand when such activities will occur; however, the particulars of each of these events will likely not be completely predictable. For the three autonomous events that involve onboard-engineered sequences of turns, thrusting, or picture taking, the ground limits AutoNav to predetermined periods of time, allowing careful budgeting of onboard time, instrument, and computational resources. Table 4-2 is a summary of the AutoNav commands, followed by a brief description of each of the AutoNav commands and associated actions.

Table 4-2. Summary of AutoNav commands.

Command Name	Description	Arguments	Usage	Time required
Nav_Do_OD	Perform Orbit Determination.	None	1/week	10–100 min
Nav_Do_TCM	Execute a TCM.	Duration	1/week	1.5–24 h
Nav_IPS_Off_Mes*	Notify Nav of a forced "engine off."	None	1/week*	1 s
Nav_Man_Plan	Perform Maneuver Planning.	None	1/week	10–200 min
Nav_Photo_Op	Perform a Nav picture-taking and processing session; edit and store data.	Duration	1/week	1.5–8 h
Nav_Reset*	Stop all Navexec state machines.	None	Seldom*	1 s
Nav_Set_IPS	Start a mission burn.	None	1/week	5 min
Nav_Start_Encntr	Start an encounter sequence.	Seq. ID	4/encounter	1 min
Nav_Update_IPS	Update the thrust vector during a mission burn.	None	2/day	1 min
Nav_Change_Mode	Change an AutoNav operating mode.	Data vectors	2/month	5 s
Nav_Data_Downlnk	Downlink a Nav file.	File ID	2/month	20 s
Nav_Data_Update	Update a Navigation file.	File ID	2/month	20 s
Nav_IPS_Press	Pressurize the main engine.	None	1/week	1–30 min
Nav_ACM_Infoturn	Optional desired pointing of the spacecraft after a Nav event.	"Turnspec"	1/week	5 s
Nav_BBC_Deadband	Optional desired deadband of the spacecraft after a Nav event.	Deadband	1/week	5 s
*Contingency or emergency back-up command				

4.1.3.1 Nav_Do_OD. Nav_Do_OD causes Nav Main to: 1) trim the OD file data arc to the predetermined length, 2) trim the History file to a corresponding length, 3) compute data residuals and partials for all data points in the data arc, 4) estimate position, velocity, and non-grav parameters for the spacecraft state for each segment of the arc, 5) repeat steps 3 and 4 iteratively until converged, 6) write these solutions to the OD file, and 7) integrate the current best-estimated spacecraft state forward to a pre-specified time (usually about a month into the future) and write this to the Spacecraft Ephemeris file.

4.1.3.2 Nav_Do_TCM. Nav_Do_TCM causes Nav Main to perform a TCM by 1) obtaining precomputed specifications for the next TCM from the Maneuver file, 2) checking that there is a TCM scheduled within a specified time (e.g., 1 hour), 3) querying the ACS for the specifications of the turn to the attitude of the burn, 4) commanding the ACS to perform the turn, 5) if the TCM is an IPS TCM, commanding IPS to thrust for the specified time at the specified thrust, or if the TCM uses the RCS, commanding the ACS to perform the specified impulsive delta-V, 6) if there is a second (e.g., vectorized) element to the TCM, performing steps 1–6 on this leg, and 7) commanding the ACS to turn the spacecraft to the terminal attitude.

4.1.3.3 Nav_IPS_Off_Mes. The ground uses this command to inform AutoNav that IPS thrust has been forced off. This will terminate the Mission Burn State Machine, if active.

4.1.3.4 Nav_Man_Plan. Nav_Man_Plan causes AutoNav to compute the propulsive plan for the next control opportunity on the Maneuver file, if any. This may be an RCS or IPS trajectory correction maneuver or an IPS mission burn.

For a mission burn, AutoNav will 1) propagate the last spacecraft state entry on the OD file to the B-plane, obtaining the current miss vector, 2) starting with a fixed number of mission burn segments, ManPlan will compute the partial derivatives of the B-plane impact position and time with respect to burn angles of each segment and the duration of the final burn, 3) make an estimate of the changes in the burn angle and last-segment duration, 4) check the estimated angle changes for violations of pointing constraint, and if a violation occurs, then reset that angle to the constraint limit, 5) make an iteration using steps 1–4, 6) if after a fixed limit of iterations, step 5 has not converged (i.e., targeting is not "close-enough"), then mission burn segments are added to the set being updated, and steps 1–6 repeated, and 7) if the solution converges, then the Maneuver file is overwritten with the updated plan; otherwise, if there is no convergence, the Maneuver file is left unchanged.

For a TCM, AutoNav will 1) propagate the last spacecraft state entry on the OD file to the epoch of the next maneuver, 2) compute the state and state partial derivatives from that epoch to the next encounter, 3) compute the required delta-V at the maneuver time, 4) repeat steps 2 and 3 iteratively until converged, 5) determine, via interaction with the ACS, whether the desired burn direction violates spacecraft constraints, 6) if so, ask the ACS to vectorize this TCM (i.e., decompose the desired—but constrained—delta-V direction into two allowed directions), and 7) via steps 2, 3, and 4, compute the delta-V associated with each vectorized leg.

In both of these cases, a new spacecraft trajectory is computed and written to the Spacecraft Ephemeris file.

4.1.3.5 Nav_Photo_Op. Nav_Photo_Op causes AutoNav to 1) cycle through its list of candidate "beacon" asteroids, taking each in turn, 2) for each, query the ACS for the turn specifications to take the MICAS boresight to that attitude, 3) before turning, determine that there is sufficient time to turn to target, take the required pictures, and turn back to the desired terminal attitude, 4) if there is sufficient time, turn the spacecraft, 5) begin taking a sequence of pictures, sending each to the AutoNav picture-processing element when complete, 6) as each picture is processed, write its reduced data (asteroid pixel, line, pointing values) to the OPNAV file, as well as edited picture elements, 7) cycle to the next asteroid target, via steps 2–5, 8) when the list of candidates is exhausted, or the available time (as communicated in the command argument list) is exhausted, commands the spacecraft to turn to the terminal attitude, and 9) filters the contents of the OPNAV file for bad data and places the results in the OD file, whereupon the OPNAV file is optionally scheduled for downlink and deletion.

4.1.3.6 Nav_Reset. Nav_Reset causes any of the three AutoNav state machines (PhotoOP, Mission Burn, or TCM) to reset to the "off" state, if they are active.

4.1.3.7 Nav_Set_IPS. Nav_Set_IPS causes the initiation of a mission burn by 1) reading the Maneuver file and determining that a mission burn begins within a specified time, 2) querying the ACS for the specifications of the turn to the burn attitude, and 3) building and starting a sequence to start at the mandated burn start time (or immediately, if the "Set" command has

occurred within a burn segment) that turns the spacecraft and commands IPS to go to a thrusting state at the appropriate throttle level and for the specified duration.

4.1.3.8 Nav_Start_Encntr. Nav_Start_Encntr causes AutoNav to build and start a sequence that in turn starts the specified sequence at the requested encounter relative time (see Section 4.1.4.1). This command is only operable while RSEN is active.

4.1.3.9 Nav_Update_IPS. During a mission burn (i.e., after a Set_IPS command), this command causes Nav to update the current burn direction according to the Maneuver file.

4.1.3.10 Nav_Change_Mode. Nav_Change_Mode updates various control-mode flags and constant settings in AutoNav. The flags and variables set are those that need to be changed frequently or due to changes in spacecraft state or mission phase. Other, more stable, parameters are kept in the parameter files.

4.1.3.11 Nav_Data_Downlnk. Nav_Data_Downlnk causes AutoNav to downlink a specified AutoNav data file. (See Section 4.1.1.)

4.1.3.12 Nav_Data_Update. Nav_Data_Update causes AutoNav to accept a specified AutoNav data file as a replacement for an existing file. The AutoNav file of filenames is updated in this process. (See Section 4.1.1.)

4.1.3.13 Nav_IPS_Press. Nav_IPS_Press causes AutoNav to command the IPS to pressurize the plena in preparation for thrusting at the throttle level determined from the Maneuver file.

4.1.3.14 Nav_ACM_Infoturn. Nav_ACM_Infoturn allows the ground to inform AutoNav what the desired ACS turn specification is for the desired terminal attitude after a PhotoOp or TCM.

4.1.3.15 Nav_BBC_Deadband. Nav_BBC_Deadband allows the ground to inform AutoNav what the desired deadband is after a PhotoOp or TCM.

4.1.4 Uncommanded AutoNav Functions

4.1.4.1 Reduced State Encounter Navigation (RSEN) and Encounter Sequence Activation. Encounter navigation activity is performed by the RSEN AutoNav subsystem. RSEN is enabled by a Nav_Change_Mode command, whereupon the most recent estimated spacecraft state and covariance are mapped to the current time. RSEN is then activated by the receipt of an APS picture. When an APS[2] picture is received, the state and covariance are mapped to the picture time by a simple linear motion propagation, the centroid of the target is located in the frame, differenced with a predict to obtain a residual, and a Kalman-filtered estimate of spacecraft position is made. The Cartesian spacecraft state is then converted into B-plane coordinates, including linearized time of flight to closest approach; the time-of-flight information is made available to other AutoNav subsystems. This process continues with subsequent pictures, with RSEN bootstrapping states from picture time to picture time (see Figure 4-2). When AutoNav receives a Nav_Start_Encntr command, wherein Nav Main is asked to start an encounter sequence at a specific time, the time of closest approach previously computed by RSEN is

[2] After the Braille encounter, and before the Borrelly encounter, RSEN was modified to use CCD images.

compared with the current time and an absolutely timed sequence is built to start the desired sequence at the appropriate time.

4.1.4.2 Non-Grav History Accumulation. AutoNav must keep a continuous record of propulsive events by RCS and IPS onboard the spacecraft for purposes of accurately integrating the flight path of the spacecraft. In this effort, AutoNav is aided by the ACS and IPS software susbsystems, which report periodically accumulated delta-V (in the case of ACS) or impulse (in the case of IPS). The reporting period varies for the ACS, since this system buffers the accumulation and only reports when a certain threshold is crossed (typically 10 mm/s). For the IPS, the reporting is every minute. AutoNav further buffers these data under parametric control, writing permanent records in EEPROM (non-volative onboard memory) when accumulated ACS delta-V or IPS vector impulse crosses internal AutoNav thresholds.

4.1.4.3 Ephemeris Services. Ephemeris Services is the highest priority AutoNav task and is required to give ephemeris information to the ACS as often as one-second intervals under rare circumstances, but nominally is queried every few minutes. The ephemeris reads the ephemeris files of the spacecraft, the beacon asteroids, and the major planets. All of these files have Chebyshev polynomial representations of the orbital states, with velocities computed. All states are in Earth-mean-equator-2000 coordinates, as are the directions on the star catalog. Ephemeris Services also provides ephemeris data to the internal AutoNav functions as well.

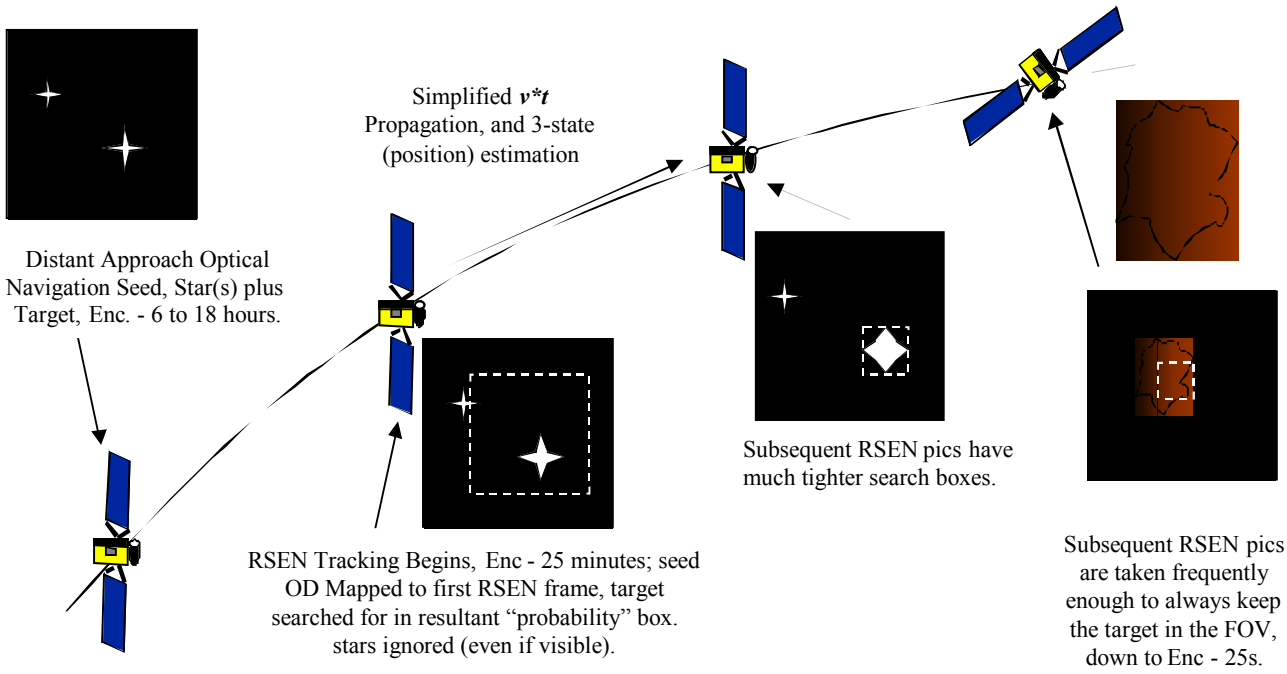

Fig. 4-2. Reduced state encounter navigation schematic functional overview.

4.2 Technology Interdependencies

4.2.1 MICAS/AutoNav Interface

The principal AutoNav dependency on other DS1 technologies is with the imaging system. For DS1, the camera is another new technology—the Miniature Imaging Camera and Spectrometer (MICAS). MICAS has two visual channels, a somewhat conventional CCD and a much smaller APS. The ability to take high-quality astrometric of small asteroids and image a bright inner solar-system target against a field of stars presents stringent requirements on a visual detector. The requirements listed in Table 4-3 were levied on MICAS, and the table indicates the level of success achieved in meeting these.

Table 4-3. Imaging system AutoNav requirements and attainment by MICAS.

	Requirement Description	Value Required	MICAS Value	Attained
1	Digitization level	≥10	12	Yes
2	Field of view	0.6 to 2.0	0.7 (CCD) 0.25 (APS)	Yes/No
3	Array size	≥512	1024 (CCD) 256 (APS)	Yes/No
4	Geometric distortion/errors	≤2 μrad	7 μrad	No
5	Device fullwell and noise	80,000 e^-/50 e^-	35,000 e^-/40 e^-	No/Yes
6	Dimmest obtainable image	12 M	9.5 M	No
7	Long-exposure capability	200 s	≤100 s	No
8	Encounter imaging	Target and 9 M	Target and 7 M	No

4.2.1.1 Overview of Camera Requirements and Attainment. Requirement 1 from Table 4-3 describes the gray levels obtainable in the instrument; 12-bit digitization, providing 4096 levels of gray, was implemented in both the CCD and APS channels, surpassing the requirement. Requirement 2, detector field of view, is met by the CCD, but not nearly by the APS. As discussed below, electronics faults in the CCD channel required AutoNav to use the APS at the Braille encounter. Additionally, light leakage and scattered light internal and external to the camera caused the effective field of view to be reduced (severely at times) in the CCD. Requirement 3 was met by the CCD but not by the APS. Requirement 4 is a complicated statement of the astrometric quality of the instrument. Factors that can affect this ability are the geometric distortion in the camera's optics, their modelability, and their temporal and/or thermal stability. Observed post-launch distortions in the MICAS optics are well over 70 μrad in extent, and due to the limiting dim magnitude of the camera, calibrations so far have been unable to improve this to better than 10 percent, or 7 μrad. Requirement 5 is a statement about the dynamic range of the instrument and the background noise. Because of the shutterless, fast-cycling readout design, the necessary range of useful signal was reduced in practice by about a factor of two from forecast, even though good noise characteristics were achieved. Requirement 6 was not achieved due to a combination of the reduced dynamic range, response-curve non-linearity, and scattered light, all to be discussed later. Requirement 7, the need to take long exposures to detect distant, beacon asteroids or the approach target, could not be achieved because of the magnitude of the scattered light problems. Requirement 8, the requirement to image the approach target with a navigation star was not met for the same reasons, substantially limiting the approach navigation strategies.

4.2.1.2 Other Camera Complications. Eight months before the launch of DS1, it was discovered that the CCD channel had a severe limitation when imaging bright objects (objects as bright as the first two expected targets). When the image of a typical asteroid brightness subtended more than 100 pixels (±50), severe charge bleed appeared in the picture due to the inability of the CCD readout to cope with the continuing photon flux. Because of this limitation, it was believed that the CCD channel would be unusable during the last few minutes of approach. Figure 4-3 shows an example of the phenomena, taken during the instrument check-out, at pre-launch. As a result of this problem, the less capable APS channel was used by AutoNav on approach. In partial compensation, the readout time required for the APS was much shorter than for the CCD—2 seconds vs. 20 seconds. At the first use of MICAS, it was apparent that there were substantial light scattering problems in and around the camera. Depending upon the Sun-relative geometry, the CCD would saturate, achieving maximum measurable charge in as little as 5 seconds of exposure. Because the original feasibility analysis of AutoNav called for exposures as long as 200 seconds, this clearly represented a reduction in capability by limiting usable geometries and targets.

Figures 4-4 and 4-5 show two examples of the scattered light effect in roughly anti-Sun and normal-to-Sun geometries. A third difficulty with the camera is a highly nonlinear response curve. See Figure 4-6 for the response curve of the MICAS APS, and the discussion of the encounter results in Section 5.2.2. The net effect of this electronics fault is for low-flux signals to be nonlinearly attenuated. This effect is much more severe in the APS, and it largely accounted for abnormally low throughput at the Braille encounter. Yet another substantial difficulty for AutoNav arose due to light-attenuating scratches in the optics chain over a substantial portion of the CCD center of field of view.

Additional problems and their workarounds are discussed in Section 5.2.1.1.

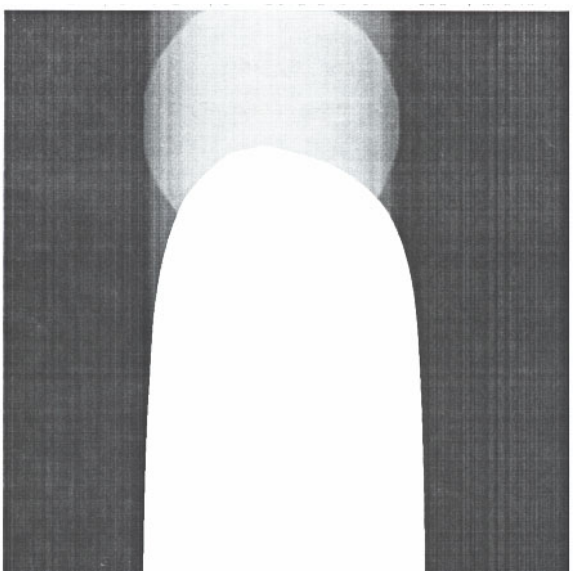

Fig. 4-3. MICAS extended bright image charge bleed.

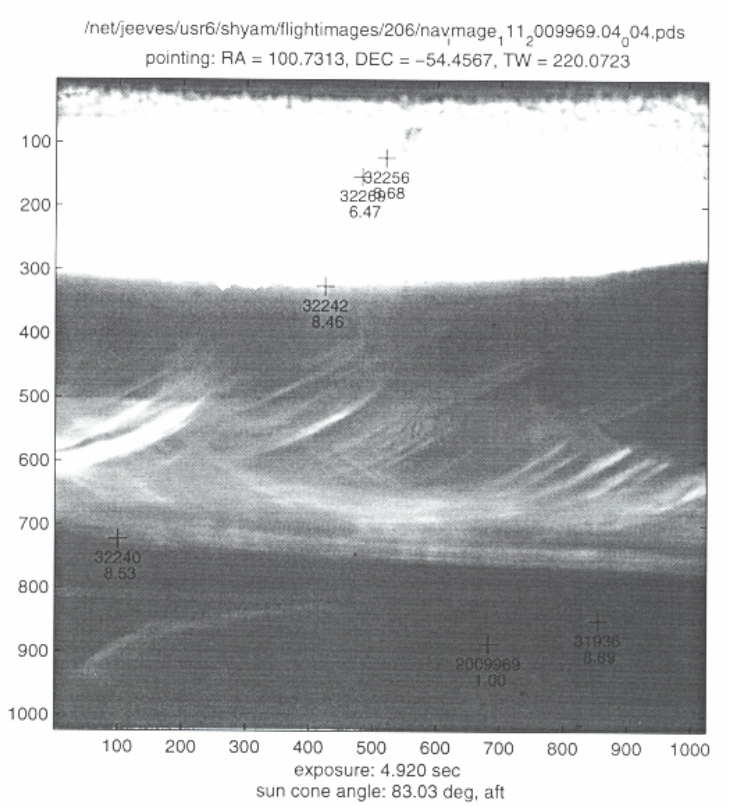

Fig. 4-4. MICAS "low solar cone angle" scattered light picture.

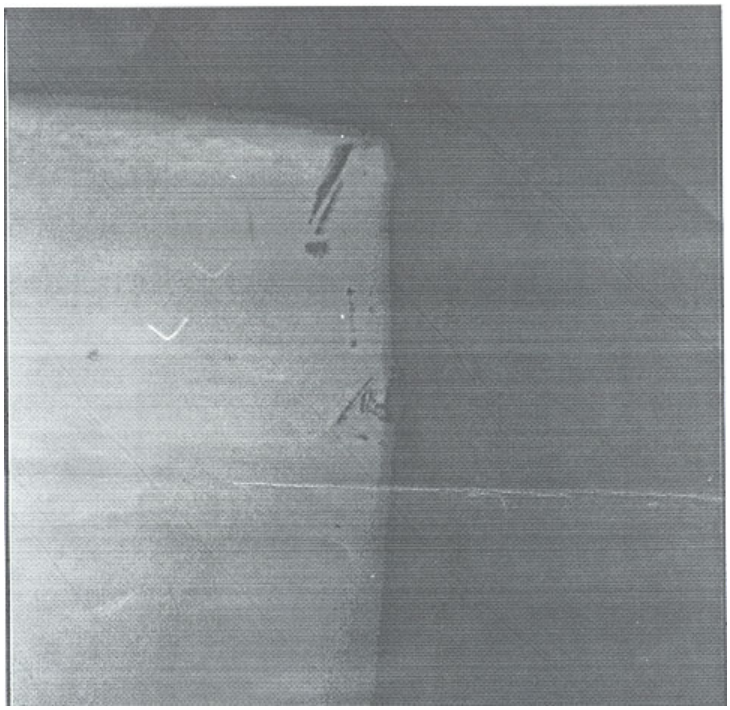

Fig. 4-5. MICAS "high solar cone angle" scattered light picture.

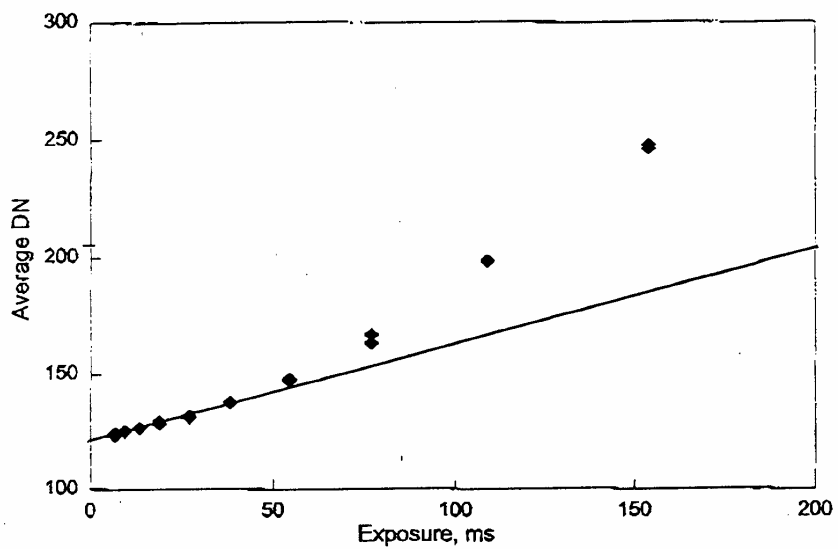

Fig. 4-6. MICAS APS channel nonlinear signal response.

4.2.1.3 MICAS Software Interactions. In addition to the MICAS hardware, AutoNav interacts with the MICAS software subsystem. The software accepts and processes requests for pictures, and provides AutoNav with important header information packaged in the picture file. The following is an example of such a header:

```
NJPL1I00PDS                      = XV_COMPATIBILITY
/*          FILE FORMAT AND LENGTH
RECORD_TYPE                      = FIXED_LENGTH
RECORD_BYTES                     = 512
FILE_RECORDS                     = 261
LABEL_RECORDS                    = 5
/*          POINTERS TO STARTING RECORDS OF MAJOR OBJECTS IN FILE
^IMAGE                           = 6
/*          ANCILLARY INFORMATION
IMAGE_NUMBER                     = 279
EXPOSURE_DURATION                =        0.013700
TARGET_NUMBER                    = 5
ONBOARD_FILENAME                 = "/micas/images/ltc300_CCD_2.pds""
IMAGE_TIME                       = 58028726.921814
SC_SUN_POSITION_VECTOR            = {109905396.260058,-129004901.095362,-
56328752.753662}
SC_SUN_VELOCITY_VECTOR             = {       19.890484,       17.517464,
7.523768}
SC_ATTITUDE_QUATERNION             = {        0.325205,        0.512832,
0.767046,       0.207087}
DETECTOR                         = "VISCCD"
IMAGE_USE                        = "SCI"
READOUT_CLOCK                    = "300KHZ"
MIN_COMPRESSION_RATIO            =     1.00
UV_VOLTAGE_LEVEL                 = 13
OBA1_TEMP                        = -123.66
OBA2_TEMP                        = -126.63
OBA3_TEMP                        = -124.74
M1_MIRROR_TEMP                   = -124.04
IR_RADIATOR_TEMP                 = -165.26
OBA_CUBE_SUPPORT_TEMP            = -124.20
IR_DETECTOR_TEMP                 = -160.21
UV_DETECTOR_TEMP                 = -5.90
ELECTRONICS_CHASSIS_TEMP         = 29.52
COVER_ACTUATOR_TEMP              = -10.85
```

```
SUBIMAGE_X                      = 132
SUBIMAGE_Y                      = 640
CLIENT_DATA                                                                 =
0x0000000000000000000000000000000000000000000000000000000000000000000000000
00000000
0
/*          DESCRIPTION OF THE OBJECTS CONTAINED IN FILE
OBJECT                          = IMAGE
 LINES                          = 256
 LINE_SAMPLES                   = 256
```

As well as taking and providing the images, the MICAS software set also compresses images with varying ratios of loss, from 1.0 (no loss) to small fractions, and consequently high compression. The software will also edit a picture to extract specified regions.

4.2.2 Attitude Control System (ACS)

AutoNav has mission-critical interfaces with the ACS. Basic spacecraft health is dependent upon AutoNav providing the ACS with the locations of the spacecraft and requested target bodies. Without this information, the spacecraft will be forced (under certain circumstances) into safing. In order to accomplish its autonomous activities, AutoNav communicates with the ACS in several ways. Though not explicitly called out as a technology demonstration of DS1, the design and implementation of the DS1 ACS contain a number of important technological advances. These include the operation of the IPS, attitude maintenance and turns with highly constrained attitudes, and autonomous turn planning for AutoNav.

4.2.2.1 Turn Planning and Execution. The ACS's Attitude Planning Expert is the service AutoNav uses to plan turns. When NavExec changes the attitude of the spacecraft, it queries the APE for the particulars of the turn between the assumed beginning attitude and the desired attitude. The APE will inform NavExec of 1) whether the turn is possible at all, 2) whether the turn violates (or nearly violates) any pointing constraints, and 3) how long the turn will take. Armed with this information, NavExec decides whether to proceed. When a turn is commanded, it is accomplished with a turn specification provided by the APE. Additional attitude information is conveyed to the ACS via updates to the IPS thrust vector ("TVC-preaim" vector), which causes the ACS to effect small turns using the engine gimbals that point the throat of the ion engine.

4.2.2.2 Mode, Turn Mode, and Deadband Changes. During the course of its autonomous work, AutoNav occasionally needs to alter the operational state of the ACS. These changes include changing from normal RCS mode to Thrust Vector Control (TVC) mode when operating the IPS is required. The mode that controls the pairs of thrusters used to turn the spacecraft must also be altered to allow for slow, deadband maintenance during picture taking. For most of the spacecraft actions that AutoNav commands, the attitude-control deadband itself must be changed to suit the activity. In addition, the ground-generated sequence must set the family of constraints that prevents areas on the spacecraft from incurring Sun-illumination before certain AutoNav events.

4.2.2.3 Queries for Current State and Delta-V Estimator. As stated earlier, the ACS periodically queries NavRT for ephemeris information. These queries always include a request for the spacecraft position and a request for the position of the body (if any) toward which the spacecraft is currently pointing. The ACS also records all propulsive activity from the RCS and

computes a net translational change in velocity (delta-V). When the value of this delta-V is greater than a predetermined value, a message containing the accumulation is sent to AutoNav, and after further buffering, these quantities are recorded in AutoNav's NonGrav History file.

4.2.2.4 Vectorization and Delta-V Requests. Because of Sun-illumination constraints and geometric constraints involving keeping the solar panels focused on the Sun, it is impossible to point the spacecraft in certain directions. If it is necessary to accomplish a TCM in one of these directions, the vector must be broken up into two components that are allowed. The APE provides a service wherein AutoNav requests a delta-V direction and APE responds with one or two allowed directions for burning the engines. Upon receipt of this information, AutoNav recomputes the magnitudes of the burn elements if it has been vectorized. When the final values of the TCM have been computed, NavExec turns the spacecraft (through interaction with the ACS) and either asks for an RCS delta-V or causes the IPS to burn for a specified time.

4.2.3 Ion Propulsion System (IPS)

AutoNav has responsibility to perform basic operation of the IPS during mission burns and TCMs that use the IPS. Additionally, the IPS reports to Nav Main the progress of any IPS thrusting. NavExec commands IPS through directives to pressurize at a given thrust level, ignite the engine, and stop and safe the engine. The IPS in turn gives reports of the accumulated impulse over a one-minute period and reports when the specified duration of the burn has been achieved. When this last message is received, Nav commands the engine to shut down. Accumulated IPS impulse is recorded in the NonGrav History file.

4.2.4 Fault Protection

One of the fundamental guidelines in the design of the AutoNav system was to minimize the possible amount of trouble that the system could cause other systems or the spacecraft overall. AutoNav to a very large degree attempts to trap all of its possible errors internally and exit the faulty function in a manner that to the external system looks normal. As a result, there were no explicit connections to the fault protection (FP) system. It was additionally felt that none of the types of internal AutoNav failures mentioned above warranted notice by FP, even in a monitoring sense. Furthermore, the general use of the sequencing system for most commanding that involved actual spacecraft actions meant that AutoNav requests for action were covered by the usual FP provided by any sequence. There is one indirect method by which FP can detect an AutoNav failure. During certain fault recovery modes when the ACS does not receive ephemeris data from AutoNav, it complains to FP, which will variously, depending upon circumstances, merely note or take the spacecraft to a higher level of fault state. As part of a safing event, FP will run scripts that set the AutoNav modes into stand-by states wherein no attempts will be made to alter EEPROM files, including the NonGrav History file.

Section 5

Performance

5.1 Radio Performance

In addition to being used to provide the DSN with trajectory updates, the radio navigation solutions were used as a baseline against which to compare the performance of the AutoNav during the cruise phase of the mission. Direct comparisions of optical solutions against selected radio solutions can be found in Section 5.2. Ideally, long-arc (several-month) baselines were used. Estimation of the trajectory using radio observations taken during a sample thrust arc follows.

5.1.1 A Sample of a Long-Term Arc Estimation

The DS1 thrusting profile of March 16 to April 27, 1999, was broken into roughly one-week segments, with an AutoNav opnav image acquisition and processing session and HGA communications pass following each thrusting period. Radio tracking data were limited to Doppler and ranging during the 10-hour HGA pass and several LGA Doppler-only tracks of about 4 hours during IPS thrusting. The IPS start was scheduled such that the first thrusting on each segment was visible in the Doppler data at the end of the HGA tracking pass.

The nominal thrust model included discrete changes in thrust (due to changing power availability as a result of solar distance changes) and direction, the latter at roughly 12-hour increments to closely approximate the Sun-relative variation of the thrust vector. In practice, autonomous battery management algorithms and ground commands changed the thrust level many more times than the 1 to 2 changes that would be typical due to power changes. Consequently, the recorded data returned from each HGA pass and the command history were required before the weekly orbit solution could be obtained. The NonGrav History file was also used to construct a delta-V model for each opnav session, consisting of an impulsive delta-V at the beginning and end of each opnav to match the total velocity and position change reported for that interval.

The orbit determination approach was generally similar to that used for IPS calibrations (see Section 5.1.2). The stochastic update rates for thrust changes was often much longer than during IPS calibrations due to long periods in which the spacecraft was not directly observed from the ground. Also, the impulsive delta-V *a priori* uncertainties were much larger, since each opnav session produced a delta-V of up to 100 mm/s (mostly from the tight opnav pointing deadbands being maintained by the unbalanced RCS Z-thrusters). The few occasions when tracking data were collected for the first few hours of IPS thrusting were used to construct an approximate model for use without this visibility, and the sequenced, commanded, and autonomous IPS thrust-level changes were included in the thrust model. Even with the significant amount of effort required to obtain orbit solutions extending through the preceding HGA pass, precise prediction to the next target (asteroid Braille) was impossible due to the unplanned thrust changes that invariably occurred. However, reasonable orbit determination accuracy was maintained within the data arc.

As mentioned in Sections 2.1 and 3.1.3, DS1 orbit determination benefited from having at least one tracking pass per week with the IPS off and relatively low line-of-sight acceleration errors from RCS activity. The HGA pass was also often split between two stations, which was quite helpful for determining geocentric declination, particularly in the several cases where the two stations were in the Goldstone and Canberra complexes, forming a North–South baseline component. Every orbit solution produced a state covariance at the end of the data arc, so the associated state estimate can be compared to solutions from following weeks, where the comparison epoch is within the data arc and presumably better determined. Each solution's propagation through the following week can be compared with a later reconstruction as well, with an assumed position uncertainty of 150 km based on 1 percent thrust magnitude and 10-μrad thrust direction errors. Figure 5-1 shows that the covariance and errors were consistent within the data arc, but that a one-week prediction had errors of over 450 km, mostly due to the unpredictable thrust-level changes.

5.1.2 IPS Acceptance Tests

5.1.2.1 IAT1. The first use of the IPS in its full operational mode was planned as an acceptance test, the IPS Acceptance Test 1 (IAT1), to assess and measure engine performance. The plans for IAT1 included three spacecraft turns to rotate DS1 by 30 degrees about its Y- and X-axes from the original attitude, defined by pointing the –Z-axis at the Earth for maximum Doppler sensitivity during thrust calibration. The intent of this activity was to obtain a three-dimensional estimate of the thrust vector direction in spacecraft coordinates, although in retrospect this would have been difficult to achieve to a useful accuracy due to thrust-level variations. Following the last turn (back to –Z to Earth), the IPS was to spend 1.5 hours at each of 6 increasing thrust levels (from 20 mN up to almost 80 mN), before being shut down either by command or autonomously as a result of overloading the solar array.

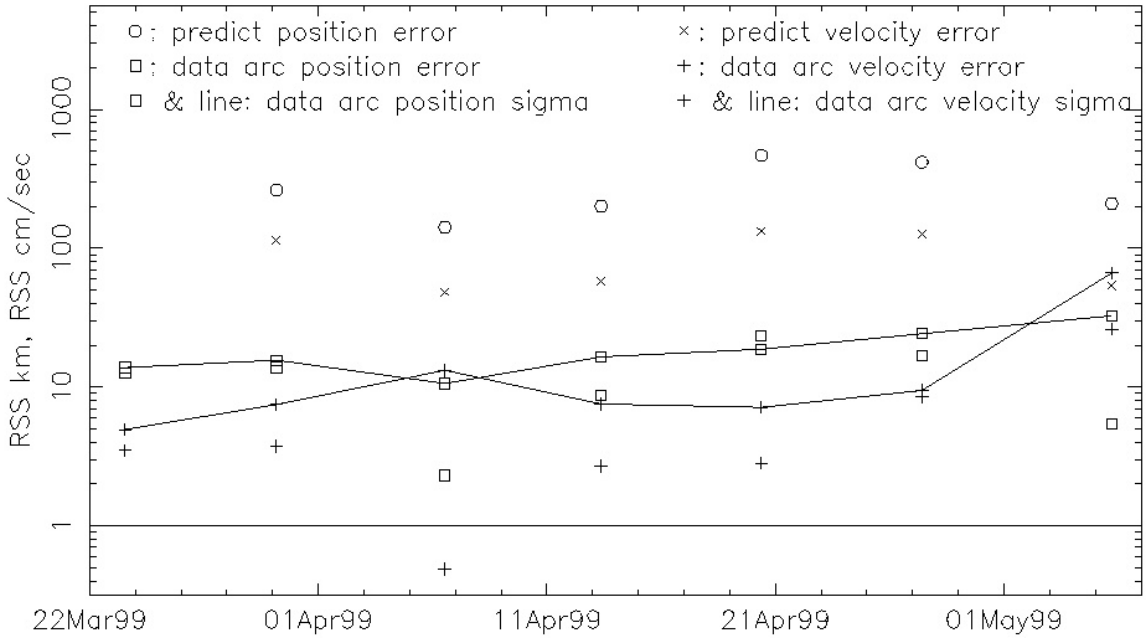

Fig. 5-1. Orbit solution consistency within the data arc and with one week's prediction.

On November 10, 1998, IAT1 proceeded nominally through the various attitude control system modes leading up to the start of IPS thrust generation. However, after 4.5 minutes, the IPS shut itself down due to the detection of a short across the accelerator grids. Continued efforts were made to restart the IPS during the remainder of the planned IAT1 sequence without success, so the IPS team began analysis and testing to isolate and correct the problem. While this effort continued for two weeks, other aspects of the spacecraft operations were tested, including optical navigation image acquisition and spacecraft safe-mode entry (twice, inadvertently). Even though IAT1 only produced 12 mm/s of IPS delta-V, along with a significant amount of RCS delta-V, the Doppler tracking data obtained by the DSN provided an exquisite measure of this result due to system performance characterized by the typical observed data noise frequently dropping to as low as 0.02 mm/s (and never much worse than 0.05 mm/s) for 10-s integration times over the first several months of the DS1 mission.

5.1.2.2 IPS First Thrust Arc Results. The next attempt to start IPS thrusting came on November 24, 1998, as the continuation of a series of IPS diagnostic tests. To the general delight of the DS1 flight team, the IPS came on and stayed on at the lowest thrust level as commanded, causing an immediate replanning effort to accomplish mission goals. The IPS thrust level was raised twice the following day before being lowered for the duration of the upcoming 4-day weekend (JPL Thanksgiving Day holidays). On November 30, 1998, four more higher thrust levels were exercised, and the IPS was left at the highest sustainable level to accumulate useful delta-V in accomplishing the asteroid encounter. The thrust attitude used for the first 10 days of IPS operation was not optimal for either trajectory change or IPS calibration, since the nominal attitude (spacecraft +X-axis to the Sun, with about 43 deg from Earth to the spacecraft –Z-axis, the IPS exhaust direction) was used to generate the IPS starting sequence without any strong expectation that IPS thrusting would continue for long at that attitude. However, the thrust direction was helpful for obtaining useful delta-V, and the thrust component in the Earth direction was not seriously degraded, allowing useful results to be obtained.

5.1.2.3 IAT2 Results. IPS Acceptance Test 2 (IAT2) occurred on May 28, 1999, with a goal of evaluating IPS performance changes after a significant amount of use, which in this case amounted to nearly 1800 hours of IPS thrusting since launch. Since the geocentric range precluded any telemetry except on the HGA, the first 5 hours were spent thrusting perpendicular to the Earth direction (a consequence of the HGA mounting geometry) before turning the spacecraft –Z-axis towards the Earth for maximum thrust observability. By the completion of the turn, the IPS pressures had reached steady-state levels, and a series of 5 thrust levels was exercised. Table 5-1 shows IAT2 estimation assumptions, which are similar to those used for IPS1. The main differences were a tighter IPS thrust stochastic *a priori* σ and the absence of a stochastic RCS delta-V scale factor. The Doppler data quality on the –Z LGA was not as good as that of the earlier IPS1 arcs, even without telemetry, due to the increased range to the Earth. But there is no evidence of IPS plasma interference with the X-band radio signal.

The IAT2 thrust estimates from the Earth-pointed period are shown in Figure 5-2. The engine throttling strategy requires that the xenon flow rate decrease slightly from the 21 mN level ingoing to 24 and 27 mN, which may be the cause of the slight downward trend in the first hour of thrusting at 27 mN. Aside from the first thrust level, 31 mN is the only thrust level for which an exact comparison exists in IPS1, and unfortunately RCS activity (most likely caused by IPS thrust dropouts, as described earlier) significantly disturbed the last half hour at this level. The final thrust level only lasted 25 minutes (as planned), due to concerns about battery discharge while operating at this power level. Thirty minutes after the sequenced IPS shutdown,

the spacecraft started a turn back to the nominal HGA Earth-pointed attitude to return recorded telemetry from IAT2. The constant thrust uncertainty for IAT2 is 0.04 mN, which is probably so small due to the relatively short duration of the thrust and the Doppler and range data on each end of the thrusting. While the in-plane component of thrust direction is well determined at 0.12 deg, the normal component is not improved from the *a priori* uncertainty of 1 deg.

5.1.2.4 Comparison of IPS Performance.

Since the total thrust estimate changes frequently with the stochastic update rates used in these solutions, average thrust levels for comparison purposes are obtained by statistically combining the thrust estimates over representative time spans and the constant thrust estimate over the entire arc. In addition to the formal uncertainty, the standard deviation of the individual estimates is computed, and whichever value is larger is adopted as the reported uncertainty. Consequently, the averaging time span is chosen to minimize the variability in the estimated thrust, even at some cost in formal uncertainty. The average thrust levels from IPS1 (formerly known as IAT1) and IAT2 are shown in Table 5-2, along with the nominal thrust (followed in the same cell by the thrust-level index from a range of 112 possible values). Since the constant thrust level uncertainty from the November 30–December 1 arc of IPS1 is 0.4 mN, as mentioned above, IPS1 values from November 25 are used up through 48 mN. The percentage difference from the pre-flight nominal is given for each average thrust. While the lowest thrust shows a half-percent increase from nominal, all other thrust levels are 1.1 to 1.5 percent lower. The IAT2 results show a consistent additional 0.5 to 1.0 percent degradation, although there are only two common thrust levels. These results suggest that future IPS use on longer missions should allow for at least a 1 percent degradation.

Table 5-1. IAT2 estimation assumptions.

Estimated Parameters			
Parameter		*A priori* σ	
Impulsive maneuver components		Z-axis: 1.0 mm/s	
		X-,Y-axis: 0.5 mm/s	
Thrust start time		1 s	
Thrust duration		1 s	
Data Weights			
Data type		Weight (1σ)	
Doppler (60 s)		0.1 mm/s	
Range		1.0 m	
Estimated and Stochastic Parameters			
Parameter	Estimated *a priori* σ	Stochastic *a priori* σ	Update Rate/Correl. Time
Spacecraft Accelerations			
Z-axis	10^{-10} km/s^2	10^{-11} km/s^2	1 hour/none
(Z, X briefly)		5×10^{-9} km/s^2	5 min/none
X-, Y-axis	10^{-12} km/s^2	5×10^{-12} km/s^2	6 hours/none
IPS Thrust			
Magnitude	5 mN	0.5 mN	2–5 min/none
Direction	1 deg	0.5 deg	1 hour/6hours
RCS Δ-V scale	Not estimated	50%	per Δ-V

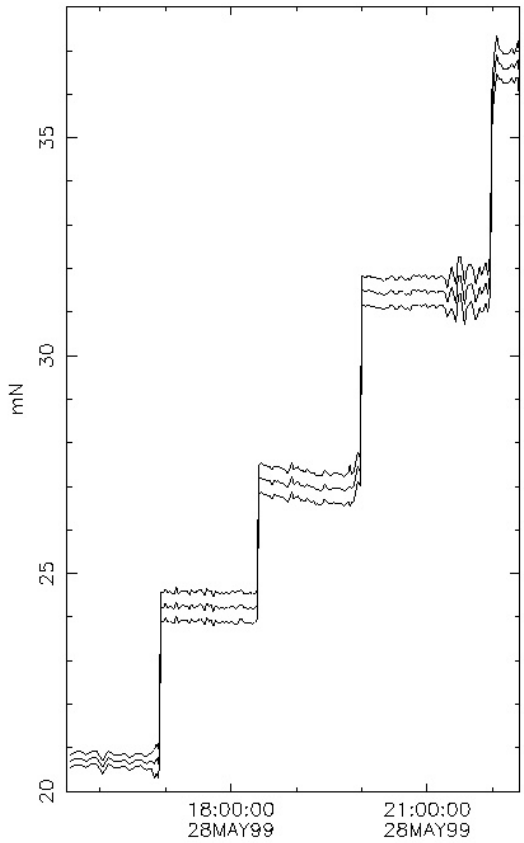

Fig. 5-2. IAT2 thrust estimates and 1-σ error corridor.

Table 5-2. IPS thrust estimates comparison.

Nominal Thrust, mN (thrust level index)	IPS1 thrust, mN (percent from nominal)	IAT2 thrust, mN (percent from nominal)
20.7 (#6)	20.797 ± 0.125 (+0.49 ± 0.60)	20.705 ± 0.082 (+0.05 ± 0.40)
24.6 (#13)		24.234 ± 0.065 (–1.29 ± 0.26)
27.5 (#20)		26.985 ± 0.073 (–1.75 ± 0.27)
32.1 (#2)	31.766 ± 0.214 (–1.10 ± 0.67)	31.460 ± 0.074 (–2.05 ± 0.23)
37.4 (#34)		36.616 ± 0.231 (–1.98 ± 0.62)
47.9 (#48)	47.298 ± 0.140 (–1.19 ± 0.29)	
63.2 (#69)	62.227 ± 0.412 (–1.49 ± 0.65)	
73.6 (#83)	72.561 ± 0.408 (–1.41 ± 0.55)	
78.4 (#90)	77.388 ± 0.449 (–1.27± 0.57)	

5.2 Optical Performance

Table 5-3, the summary of the AutoNav validation plan and success, gives a succinct summary of all of the validation events undertaken. Where applicable, quantitative goals and achievement levels are listed. In general, there is a range of achievement in these values, and where this is so, best and worst values are noted. In the body of this summary, especially Section 5.2.1.1, the history and conditions of these variously good and bad results are discussed at length.

Table 5-3. AutoNav technology validation key point summary.

A	B: Technology Validation Item Description	C	D	E	F	G	H	I	J
1	Provision of Ephemeris Services	~10^5	~10^5	~10^5	0	≤0.1 km	Req'd	<<0.1 km	<<0.1 km
2	Opnav PhotoOp Process	~40	47	46	1				
2a	Picture Planning	~40	47	47	0				
2b	ACS/APE Interaction & Turn Planning	~40	47	47	0				
2c	Mini-Sequence Picture/Turn/Fault Execution	~40	47	47	0				
3	Image Data Handling and Downlink	~40	47	47	0				
4	OpNav Data Accumulation, Handling, Downlink	~40	47	44	3				
5	Image Processing (RSS ensemble statistics)	~1200	~1500	~500	0	≤0.25 px	Desir'd	≤0.40 px	1.5 px
6	Orbit Determination (Accuracy within Data Arc)	~32	34	34	0	≤250 km, 1 m/s	Req'd	≤150 km, 0.2 m/s	10000 km, 7 m/s
7	Generation of Onboard Ephemeris, and Downlink	~32	34	34	0	0.1–1 km	Req'd	0.1 km	1 km
8	Trajectory Control and Maneuver Planning	~20	12	12	0				
8a	IPS Mission Burn Updates (Convergence Criteria)	~12	6	6	0	≤1 km	Desir'd	≤1 km	≤1 km
8b	IPS and RCS Maneuver Computations (do.)	~8	5	5	0	≤1 km	Desir'd	≤1 km	≤1 km
8c	TCM Execution, and Delivery (Final TCM and Accuracy – Position and Velocity)	8(2)	5(1)	5(1)	0	(≤2.5 km, 0.25 m/s)	(Req'd)	(≤1.5 km, 0.18 m/s)	(≤1.5 km, 0.18 m/s)
9	Execution of Mission Burns	~12	7	7	0				
10	Encounter Image and OD Operations (RSEN)	2	2	1	0				
10a	Image Processing, and Data Reduction	~80	~80	~40	1				
10b	Ephemeris Generation and Delivery	~80	~80	~40	0	≤0.5 km	Req'd	≤0.5 km	15 km
11	Encounter: Initiation of Encounter Sequences	8	8	8	0	≤5 s	Desir'd	≤5 s	≤15 s

Legend—A: Item Number, B: Item Description, C: No. Planned In-Flight Executions, D: No. Actual In-Flight Executions, E: No. Successes In-Flight, F: No. Failures In-Flight (due to AutoNav Fault and/or Misuse), G: Quantitative Goal Value (If Applicable), H: Required/Desired Quantitative Value, I: Best Value Achieved, J: Worst Value Achieved

5.2.1 Cruise Operations

During cruise operations, a typical AutoNav activity included a PhotoOp session, an OD attempt, and a ManPlan attempt. During a PhotoOp, the onboard navigator autonomously

collected optical observations. During OD, these observations were used to update the onboard knowledge of the current spacecraft orbit. During ManPlan, the navigator attempted to compute any necessary updates to future mission burns or planned TCMs. The following subsections describe the progress towards validation of each of these activities, including updates made to the software during the light software uploads in January and May 1999.

5.2.1.1 Image Processing Results. One key requirement for successful optical navigation is the presence of a well-performing and well-calibrated camera. For DS1, this was found not to be the case soon after launch. See Section 4.2.1.2 for a description of the effects of scattered light within the camera images. This section describes the effects of this camera behavior on the optical navigation efforts.

On November 6, 1998, the first post-launch pictures were taken with MICAS. Ground examination of these images showed serious anomalous behavior, which was later identified as significant scattered light leakage into the instrument. On December 21, 1998, the first PhotoOp attempt in which all logistical operations were successfully exercised was performed. Unfortunately, none of the pictures could be processed due to the effects of scattered light. On January 6, 1999, the onboard image-processing software was reconfigured in an attempt to compensate for the image quality. The following PhotoOp (performed on January 7) was again unsuccessful in processing any of the images.

In January 1999, an effort was undertaken to overhaul the image-processing algorithms in an attempt to cope with the severe scattered-light infiltration into the camera. Images taken during PhotoOps performed on January 18, 20, 26, and February 1 were downlinked in support of these efforts. Using the launch version of the AutoNav software and extensive parameter manipulation, only the very brightest asteroids and stars (brighter than 8.5 M) could be processed. Use of the images in this fashion allowed the ground navigators to define and test alternative image-processing software. The software changes that resulted from these efforts were uploaded to the spacecraft on February 8 as a part of a nominal software upload.

On February 18, the first PhotoOp using this new software was attempted. Unfortunately, due to a ground oversight regarding the parameter configuration, only one picture out of 30 was processed successfully. The parameters were corrected in time for the next PhotoOp, scheduled for February 22. Of 32 pictures on four lines of sight (LOS), six succeeded—three each on two lines. The next PhotoOp, on March 1, resulted in the successful processing of 13 out of 30 images. These results, while not up to pre-launch expectations, were enough to allow Orbit Determination to be performed successfully (see Section 5.2.1.2).

The MICAS and Navigation Teams undertook an extensive calibration campaign in early March to, among other reasons, attempt to characterize the scattered-light and light-leakage problems. The spacecraft imaged a pair of star clusters for purposes of calibrating the geometric "flatness" of the camera field, and these pictures revealed that there were severe distortions, up to 5 pixels in size and of unusual character. Pre-launch calibrations had indicated less than 1 pixel of relatively benign (i.e., readily calibratable) distortion in the field. With the images taken to characterize the scattered light, a quantitative analysis was made of the resulting increased noise in images, which was substantial and damaging to the navigation algorithms.

In order to cope with the camera geometric distortions, work began on a new distortion model for the flight software, incorporating a 6th order Legendre Polynomial model. To cope with the high levels of scattered light, algorithms for taking and differencing a background

picture were devised and implemented. Calibration suite images of Mars pictures indicated that the approach target (1992KD) would be very bright. Analyses of these frames indicated a nonlinearity in the CCD response, which attenuated weak signals. This nonliearity had been suspected from the earlier AutoNav frames. The result of this analysis indicated that only the brightest asteroids and stars could be processed by AutoNav. This fact required a change in strategy for picture planning. The original plan was to look at any time at a particular reasonably bright (<10.5 M) asteroid, and with the expected performance of the camera, acquire in general 2–4 10th-magnitude stars, more than sufficient for navigation purposes. However, the actual performance of the camera-limited suite of "good" asteroids was diminished by 75 percent, and the useable stars were those 9th magnitude or brighter. Consequently, far fewer asteroid or stellar targets were now available, and the Picture Planning file (Section 4.1.1.11) had to be carefully primed to allow AutoNav an opportunity to image these.

Analysis of the picture processing continued through April and May, and plans were made for further enhancements to the image processing algorithms. Substantial improvements were seen with ground processing using Legendre Polynomial corrections to the asteroid observations and using pre-processed pictures. The pre-processing entailed taking a background picture with each line of sight, and then differencing that picture from all pictures in this set. The background picture was offset slightly (approximately 200 pixels) from the navigation pictures to prevent damage to the targets. These two algorithmic changes were factored into the June flight software (FSW) load.

From June 1–9, 1999, the latest software set was uploaded to the spacecraft. This included final adaptations to the MICAS problems for cruise, including the Legendre Polynomial model of geometric distortions, and picture differencing to further reduce problems associated with scattered light. Over the next 2 months, these new elements were validated in cruise AutoNav operations.

The first PhotoOp performed with the software was unsuccessful due to the presence of an "un-updated" parameter file, which caused the image processing to work in pre-launch fashion based on hard-wired software defaults. After rectification, the next PhotoOps on May 16 and 20 resulted in 19 (of 36) and 20 (of 36) images being processed successfully.

Aside from an error found in the Legendre Polynomial coefficients (which was rectified on May 23), the onboard navigator continued to obtain optical observations at a relatively successful level:

- June 23, 1999, PhotoOp: Two asteroids were available. Sixteen pictures of each were taken and 14 and 11 of these processed successfully, respectively.
- June 29, 1999, PhotoOp: Of two available asteroids, only one processed successfully, with 12 of 16 pictures processing normally.
- July 2, 1999, PhotoOp: 28 of 36 pictures processed successfully.
- July 4, 1999, PhotoOp: 22 of 36 pictures processed normally.
- July 6, 1999, PhotoOp: 27 of 32 pictures processed normally

5.2.1.2 Cruise OD Results and How They Compared Against Radio OD. While efforts were under way to assess the performance of the image processor whose output was critical to the OD performance (see Section 5.2.1.1), the abilities of the AutoNav orbit determination software were also evaluated and improved during the course of the primary mission.

Due to the initial problems with image processing, a true OD was not performed until February 19, 1999. This OD was based on images taken during a PhotoOp on February 22 and incorporated ground-processed optical data from images taken during the PhotoOps on January 7, 20, and 26, and February 1. The 6 AutoNav-processed data added to the 36 uplinked data produced the first viable onboard autonomous OD, which was in error from the ground-determined state by about 4000 km and 2 m/s. In the interest of assessing the OD before allowing the onboard spacecraft operations to make use of it, this solution was intentionally discarded. After the OD performance was deemed to be sufficiently accurate for cruise, an update to the AutoNav control modes was made on February 27, allowing the navigator to preserve the OD results and replace the onboard ephemeris, effectively putting the spacecraft under AutoNav control after the next OD operation.

On March 1, the OD performed following the PhotoOp included the 13 successfully processed images in a data arc spanning January 5 to March 1. OD results were within 5000 km and 2 m/s of radio-navigation-determined spacecraft position. This solution was saved onboard in the form of a 60-day spacecraft ephemeris, at which point DS1 was maintaining onboard orbital knowledge based solely on optical data.

With DS1 autonomously computing its course, March activities began a period of 10 weeks of normal operations, which included weekly PhotoOp/OD/ManPlan sequences and periods of mission burns (periods of precomputed deterministic IPS thrusting). This period of regular data and high-rate downlink capability offered a good opportunity to further analyze and debug AutoNav operations. The results of the first OD attempts, even based on limited data, revealed that the optical data residuals were larger than expected. Figure 5-3 shows pre- and post-fit residuals for a solution performed onboard in this investigation period. Residuals larger than one pixel, with biases (in some cases) of several pixels, were much higher than expected. Calibration of the camera pre-launch indicated that measurements accurate to about one pixel should be obtainable without re-calibration, but apparently larger distortions developed after launch. Furthermore, measurements of stellar photometry led to speculation of strong nonlinearity in the CCD channel at low-flux levels. Necessarily, a thorough calibration of MICAS was called for, and this was scheduled for March 5. Two star clusters were chosen, one with a dense distribution of moderate to dim stars, another with a few bright stars to aid in both geometric and photometric calibration. Additionally, the MICAS team scheduled a set of calibration frames on March 11.

Between March 1 and June 1, assessment of the OD process continued. The performance was expected to be marginal until the corrections from an improved distortion model could be included in the software. During this time, two other modeling problems were identified and resolved. In mid-March, the OD performance degraded to nearly 10,000 km. In this time frame, it was realized that the RCS nongravitational modeling onboard was severely compromised due to large drops in hydrazine pressure since launch. This factor of 2 drop resulted in an approximately equal drop in specific impulse of the attitude thrusters, and thus in the modeled values of accumulated delta-V sent to AutoNav. Nevertheless, use or nonuse of this part of the model made no appreciable change in the OD performance. By late March, the OD quality degraded further to 13,000 km. However, the velocity measurements were accurate to about 1.5 m/s. This accuracy of velocity determination was inconsistent with the poor position determination, indicating that systematic biases were being observed in the astrometry. It was determined at this time that the largest share of this bias was due to an inconsistency in a model

describing the *a priori* pointing biases of the camera. These parameters were changed onboard in a subsequent parameter file upload.

By mid-April (after coping with a dearth of bright asteroids and a corresponding weakening of geometry and OD performance), the OD accuracy improved to 6,000-km position and 4-m/s velocity due to the correction of the camera-pointing *a priori* model (see Figure 5-4). By late April, this had improved to 4,000 km and 4 m/s. By the end of April and into early May, OD quality improved to 2,000 km, as the amount of corrupted data from the pointing angle *a priori* was systematically trimmed from the OD file. Velocity errors rose slightly to 4.7 m/s. Through the remainder of May, with image processing over 75 percent successful overall, OD performance improved to 1700 km and 2 m/s. This improvement was due to tuned image-processing parameters (more discrimination of image strength), the use of only strong asteroids and stars, good geometry of asteroids, and a dense late data set (see Figure 5-5).

5.2.1.2.1 OD Accuracy During Late Cruise. Following the upload of new flight software in May of 1999, the OD performance began to show marked improvement. This new flight software included an improved camera geometric distortion model, which was expected to greatly increase observation quality, and a background image filter mechanism for reducing the effects of scattered light, which was expected to greatly increase observation quantity. While these changes were aimed directly at improving the image processing (see Section 5.2.1.1), their effect was noted in the improved OD results.

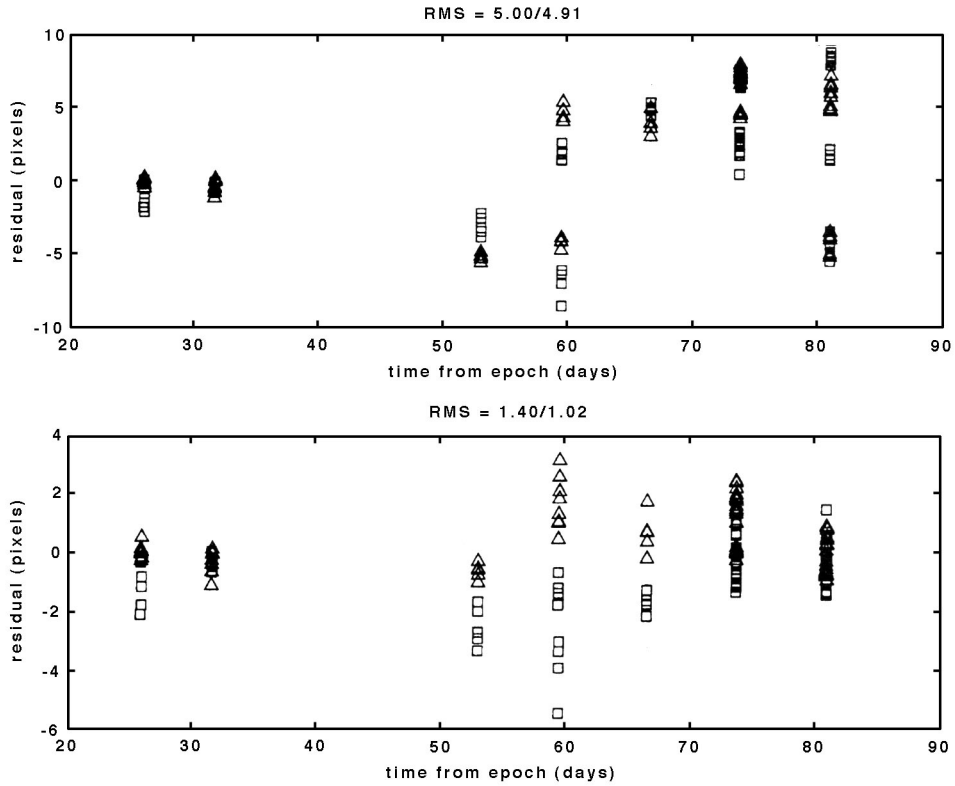

Fig. 5-3. Pre-(upper) and post-(lower) fit residuals from the March 22, 1999, optical solution.

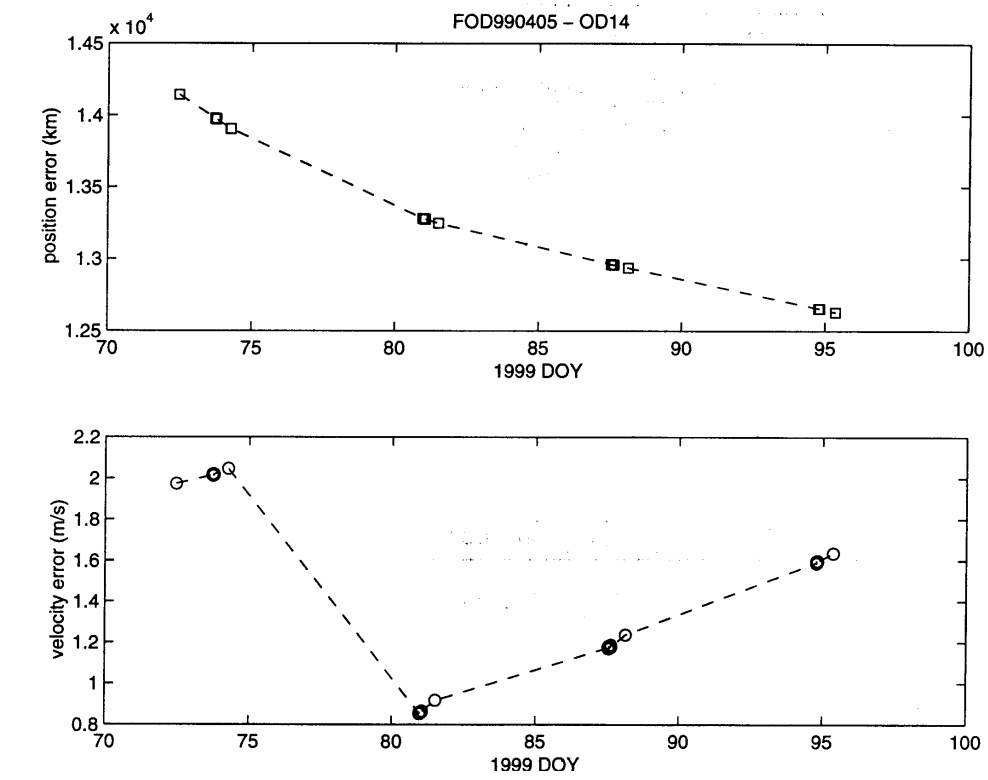

Fig. 5-4. Flight vs. ground orbit determination, April 5, 1999.

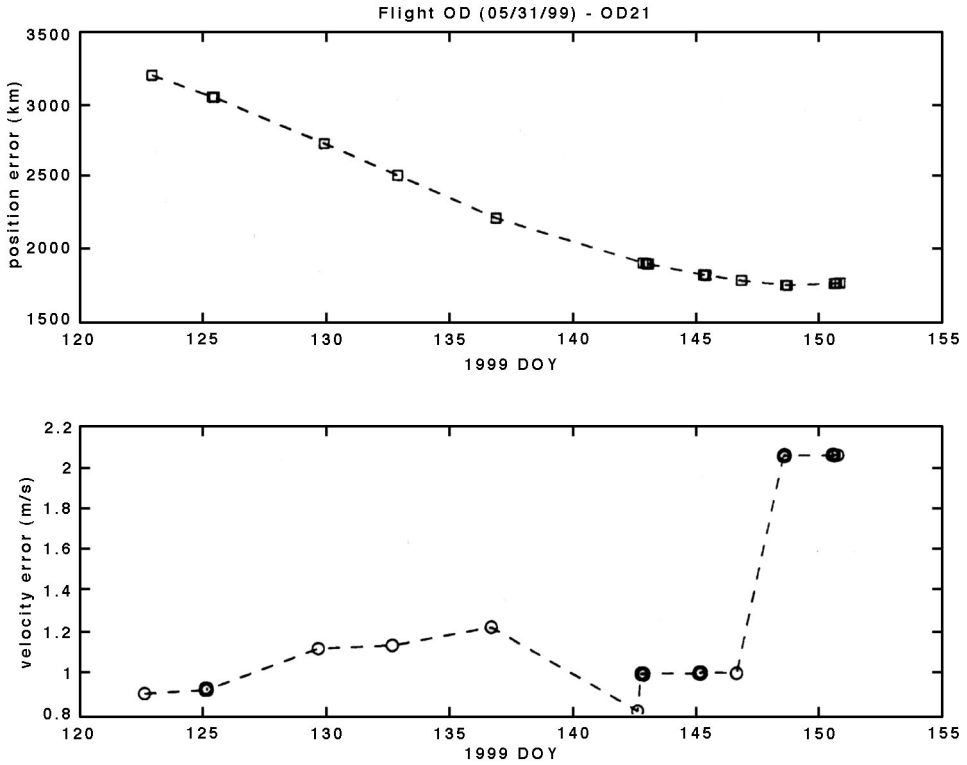

Fig. 5-5. Flight vs. ground orbit determination, May 31, 1999.

Initial attempts to use this new software were not without surprises, however. The OD solutions following the PhotoOps of June 16 and 20 degraded alarmingly to about 3500 km and 2130 km and 1.7 m/s and 0.9 m/s, respectively, from a previous position error of only 2000 km.

It was discovered that *ground* processing of the new Legendre Polynomial distortion model had been in error and that uploaded calibrated data older than June 6 were erroneous. A new OD file was prepared for uplink, with corrected calibrations, and was uplinked in late June. For OD solutions that took advantage of the fixed calibrations, the improvement was immediately noticeable. The ODs of June 29 and July 2 showed indications of the hoped-for improvement. Agreement with radio OD at this time improved to 662 km and 0.58 m/s and 904 km and 0.3 m/s, respectively.

5.2.1.2.2 OD Accuracy During Approach to Braille. In mid- to late July, as DS1 approached Braille, there remained five more OD solutions based on non-Braille objects.

The OD solutions on July 16 and 18 had an agreement with radio OD of 658 km and 0.34 m/s and 669 km and 0.32 m/s. The second of these OD solutions contained 13 Mars observations. The Mars processing invoked a previously unused mode of processing images, wherein extended bodies (Mars being about 5 pixels across) were "brightness centroided" and then that position was corrected for phase. In the OD process, these pictures were highly de-weighted (to 5 pixels as opposed to 2 for asteroids). The effect of adding Mars to the second solution did not change the solution much as a result of this de-weighting. Post-processing on the ground revealed that even with stronger weighting, Mars did not substantially improve the match between the ground radio solutions and flight, leaving unresolved the source of the observed biases, of several hundred kilometers. It was (and is currently) believed that these biases were due to a combination of residual geometric calibration defects and possibly ephemeris errors. Pre-launch, it was expected that the geometric calibration could be made to 0.1 pixel, but the insensitivity of the camera (inability to acquire dim stars) precluded this. The ephemeris errors, expected to be in the neighborhood of 100–200 km, were running somewhat larger, perhaps 400 km, as would be observed at Braille (1992KD) (see Section 5.2.2).

The Mars observations were repeated on July 19 and July 20, with onboard optical OD results maintaining a tight velocity agreement with the radio solutions (approximately 0.25 m/s), while the position errors lingered at several hundred kilometers. These Mars observations (as with the earlier Mars observations) offered unique viewing of Mars against a very bright star. Nevertheless, the substantial challenge in processing the Mars images prevented pushing the quality of the OD past the limiting effects discussed above.

In order to help compensate for camera deficiencies (largely believed to be associated with the geometric calibration), an OD file with spacecraft-acquired optical data was put onboard on July 21. These data had been "scrubbed" to remove observations that were only marginally good. With the limited data set available to the ground planners, it was impossible to set low-pass residual thresholds to a discriminating enough level to accomplish this editing onboard. In extensive ground processing and limited spacecraft processing, these scrubbed data sets were regularly achieving radio/flight OD agreements of better than 300 km and 0.25 m/s (see Figure 5-6). While this was not the final validation of AutoNav, it was quite close to achieving the necessary position agreement and more than exceeded the necessary velocity agreement. See Section 5.2.2 for a discussion on orbit determination during the encounter with asteroid Braille (1992 KD).

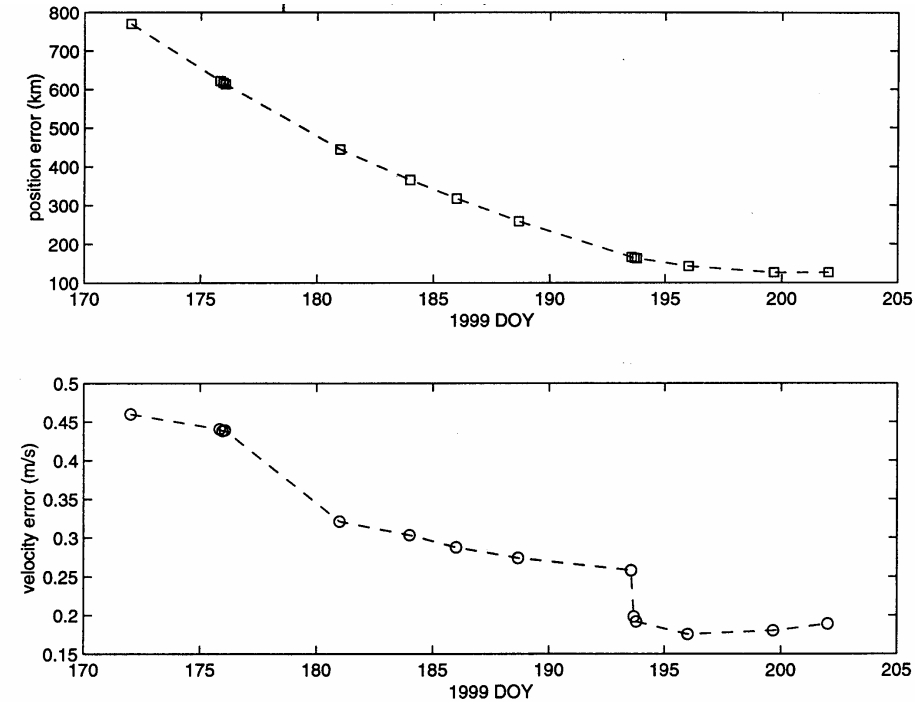

Fig. 5-6. Flight OD vs. ground OD #37, July 21, 1999.

5.2.1.3 Maneuver Execution and Planning Results. The final cruise capability of AutoNav to be validated was the autonomous execution of mission burns using the IPS, modifying the IPS thrusting profile, and planning and executing trajectory correction maneuvers using either the IPS or the RCS.

5.2.1.3.1 Maneuver Planning During Cruise. The first operation of the IPS under control of AutoNav started on December 18. Operation of the IPS in this fashion involved commanding the ACS to turn the spacecraft to the proper orientation and commanding the engine to start up at a specified throttle level. Following this, AutoNav made frequent updates to the thrust vector pointing over the next several days. The execution of the first week of this mission burn completed a few days later. On December 22, the mission burn was restarted and continued throughout the upcoming holiday season. AutoNav's commanding of the IPS in this fashion continued successfully throughout the primary mission.

During cruise operations, IPS control required periodic updates to the IPS thrusting profile. Updating the profile occurred during ManPlan activities, which were routinely scheduled to occur following each PhotoOp and subsequent OD (see Sections 5.2.1.1 and 5.2.1.2). The first ManPlan activity was commanded to occur following the OD of February 22. It resulted in correctly deciding that no upcoming maneuvers needed to be modified.

With the commencement of the second mission burn on March 16, a ManPlan was executed in the presence of a control opportunity. ManPlan correctly assessed that the OD uncertainties (the OD filter formal errors; see Section 5.2.1.2) mapped to the 1992KD encounter were larger than the mapped position errors. Consequently, thrusting on the nominal plan continued.

In the interest of assessing the performance of the ManPlan maneuver computations, a change was made to the AutoNav parameters such that the mission burn profile would be updated regardless of the formal uncertainties of the OD solution. The effect of this change was noted on April 12, during the maneuver plan window that followed the OD on that day. Based on this OD (which was within 6000 km and 4 m/s to ground OD), a thrust profile update was successfully computed and deemed adequate to use for the upcoming mission burn. On April 26, another maneuver planning attempt followed the OD of that day. While the OD solution quality had improved (see Section 5.2.1.2.1), the ManPlan execution for the associated mission burn was unsuccessful. This was due to the shortness of the remaining burn arc and the fact that ManPlan was forced to compute statistically insignificant changes. As a result, the nominal plan was reverted to onboard.

5.2.1.3.2 Maneuver Planning During Approach to Braille. Maneuver planning continued during the cruise phase of the mission, as it was used to modify the thrusting profile to arrive at asteroid Braille on schedule. As the encounter neared, improvement to the image processing and OD software (see Sections 5.2.1 and 5.2.2) allowed for the calculation and execution of TCMs, which were needed to successfully target the final flyby of the asteroid.

The first of these TCMs was scheduled for May 14, and a successful plan for this TCM was computed on May 10. The decision criterion used (to decide whether or not to utilize the onboard computed maneuver parameters) was that it was necessary for AutoNav to reduce the distance remaining to the target (projected into the B-plane) by at least half. In this case, the criterion was satisfied. This was computed to be a 1.5-m/s IPS TCM, vectorized along two legs, to correct 830 km in the 1992KD B-plane and 58 seconds time of flight (or 870 km). No problems were encountered during the successful execution of this TCM on May 14.

On July 6, the IPS TCM for Encounter minus 20 (E-20) days was computed. Figure 5-7 shows an assortment of OD solutions from AutoNav onboard, AutoNav mirrored operations on the ground, and radio navigation. Within this complex figure, it can be discerned that the AutoNav solution of July 6 created a TCM solution (when measured against the radio solution) that would not meet the acceptance criteria for an autonomous TCM—namely reducing the B-plane error by one-half. This solution would have met the criteria, had a small change in nongrav modeling procedure not inadvertantly been changed to be lacking a forecast of delta-V's associated with PhotoOps. This change caused a 400-km discrepancy in the solution (well within the formal uncertainties, as shown), which was enough to cause the maneuver solution not to meet the criteria. Since several upcoming TCM opportunities existed, it was decided to cancel the –20 day TCM.

In preparation for the planning and execution of the next TCM, scheduled for E-5 days (encounter minus 5 days), two files were placed onboard the spacecraft. On July 19, the final best-ground-determined Braille ephemeris was loaded onboard the spacecraft, representing the observing efforts of about a dozen astronomers over 18 months, and incorporating observations less than 2 weeks old. It was believed that this ephemeris was accurate to about 150 km (1-sigma). Also, a maneuver file was placed onboard with a TCM design based on the radio data (see Figure 5-8). If, after the July 22 PhotoOp, it was decided that the onboard planned TCM design was inadequate (the decision criterion was to reduce the net deflection from target by one-half), then this file would have been made the primary maneuver file.

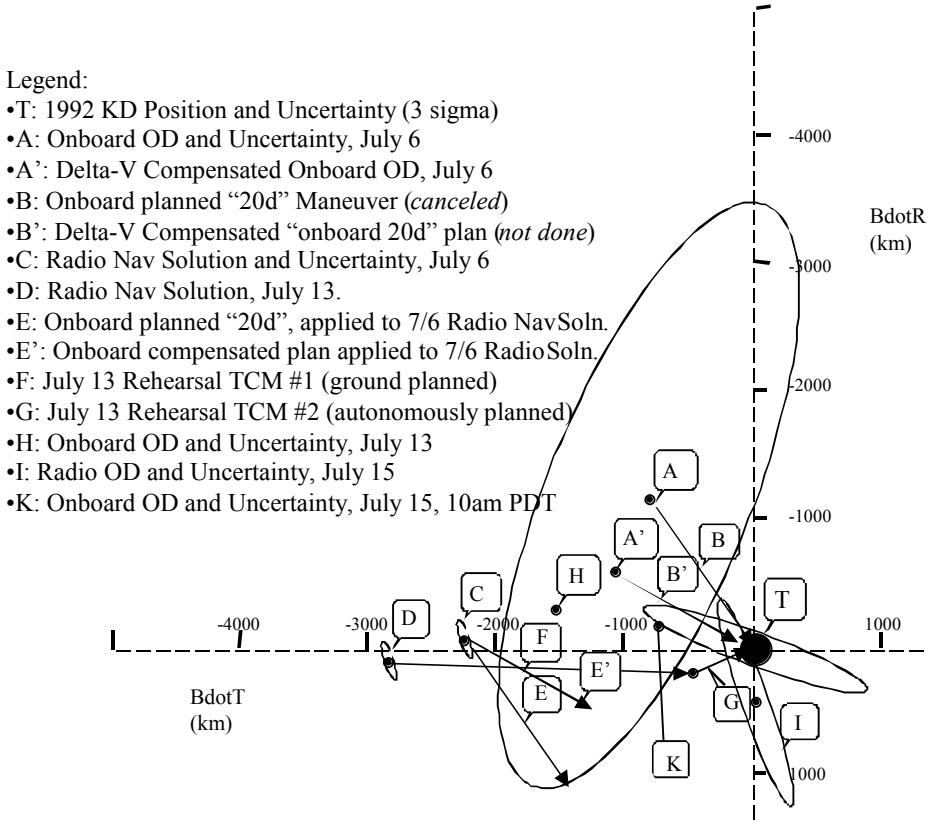

Fig. 5-7. Current B-plane target conditions at the –20 and –10 day TCMs: decision data from July 15, 1999.

On July 22, before what was to have been the ManPlan to compute the E-5 day TCM, a problem occurred—the source of which was never identified—that caused one or more of the pictures (in this case, an image of Mars) to be off-pointed (see Section 5.2.1). This in turn tripped a latent AutoNav software bug, which caused the erroneous writing of large blocks of data into the OPNAV file. This effectively filled the flight software file's file system. The OPNAV file was unreadable by AutoNav, and so the OD function failed, reverting to the *a priori* OD file. This solution was within 250 km of the radio solution "at epoch" (e.g., on July 21) and mapped to an error of 400 km in the Braille B-plane (see Figure 5-9). This solution did meet the acceptance criteria for the onboard TCM design, but only barely. However, because there was an associated anomaly with the PhotoOp and OD, it was decided to revert to the ground design. This was accomplished with a simple Nav_Data_Update command (see Sections 4.1.3.12 and 4.1.1) to point AutoNav to the already onboard maneuver file. This anomaly had the beneficial effect of alerting the AutoNav team to this bug, which posed a threat to the close-approach sequences. The Picplan File (picture planning file) was changed at the next opportunity to ensure that extended-image picture processing would not be used in any of the subsequent PhotoOps, as was then planned for those within 5 hours. With this picture-taking mode disabled, it was believed that AutoNav would receive insufficient improvement in position from the early approach pictures to warrant the –3-hour TCM. Consequently, the sequence for this TCM was altered, and the Nav_Do_TCM call was replaced with a simple turn to a Braille pointing attitude. On July 23, the TCM in question executed normally. Figure 5-8 shows the effect of the TCM, approximately 500 km in the "B • R" direction.

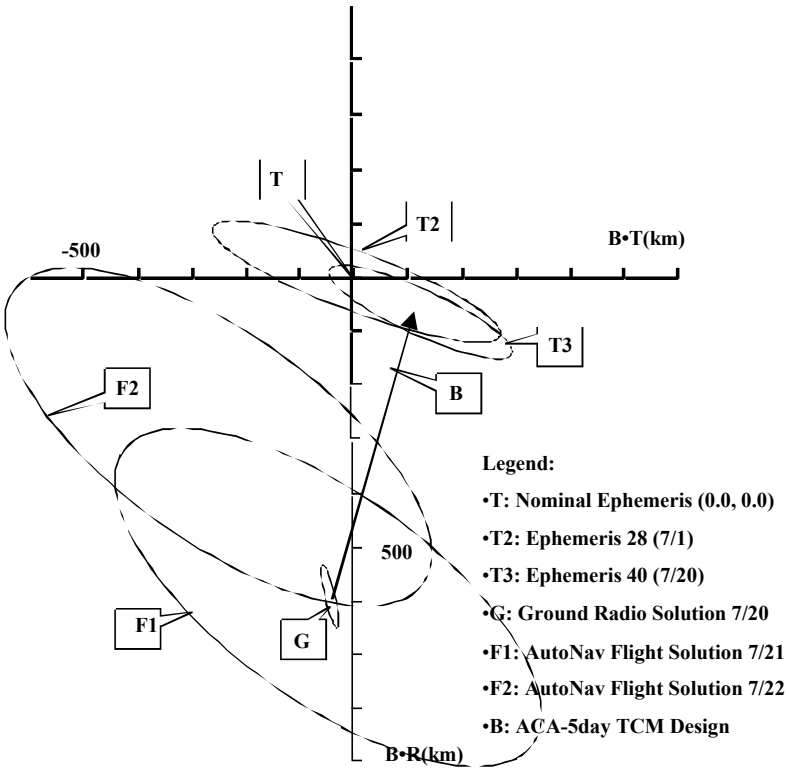

Fig. 5-8. E-5-day TCM solutions.

All further TCMs (E-2 day, E-1 day, E-18 hr, and E-6 hr) were planned and executed using conventional non-autonomous means, due to the fact that the very low signal strength of the Braille images was inadequate to autonomously process. These results are discussed in Section 5.2.2.2.

5.2.2 Encounter Approach, Preparation, and Results

Preparations for encounter and for the encounter rehearsal began early in 1999, but focused only on the last 90 minutes of operations until March, when the activities of the last 6 hours before closest approach were planned. By early July, the details of the last two days had been planned. Table 5-4 summarizes the navigation-related activities and durations of the last two days.

5.2.2.1 The Encounter Rehearsal. The encounter rehearsal, originally scheduled for June 25, involved an extensive series of practice runs on the testbed and set-up activity on the spacecraft. In order to accomplish these, rehearsal files had to be created, including spacecraft ephemeris, simulated body ephemeris, a target star catalog, and tailored parameter files. These data created a simulated universe in which the spacecraft found itself upon initialization of the rehearsal. Within this universe, the spacecraft "saw" (through the onboard test simulator) modified images, the phantom approach target (dubbed "Spoof"), and computed its position relative to Spoof, adjusting course correspondingly. It was desired (and necessary) to use the rehearsal as the first execution of an RCS TCM. It was further desired to use the approach TCM to Spoof to correct the actual approach asymptote to 1992KD. The rehearsal maneuver file was tailored to make the

first of the rehearsal TCMs (that is, for the rehearsal only) deterministic. This TCM was a ground-designed event that removed much of the then existing residual in the B-plane. At the same time, sufficient residual needed to be left for the second of the two rehearsal TCMs to be a substantive test, and not to endanger the 1992KD encounter if it misfired in any way. The files for the rehearsal were uploaded to the spacecraft on June 23, while ground tests in the testbed continued. The results from these tests were good from an AutoNav standpoint, with Nav tracking the target spoofed to within 30 seconds of closest approach. However, there was substantial uncertainty about other subsystems, so the June 25 onboard rehearsal was canceled and rescheduled for July 13. Aside from the requirement that all of the encounter rehearsal-specific files be regenerated, this also resulted in the loss of any opportunity to update the flight software if problems during the rehearsal had been encountered.

Table 5-4. Navigation encounter activities.

Encounter Relative Event Time	Duration	Activity	Sequence No.
-2d 3hr	180m	RCS TCM ("Minus 2 Day")	AN300
-2d 0hr	210m	PhotoOp/OD/ManPlan	AN301
-1d 21hr	240m	High Gain on Earth Telecom Track	
-1d 17hr	210m	PhotoOp/OD/ManPlan	AN301
-1d 14hr	240m	High Gain on Earth Telecom Track	
-1d 10hr	210m	PhotoOp/OD/ManPlan (OD and Maneuver Planning for –1d TCM)	AN301
-1d 3hr	180m	RCS TCM ("Minus 1 Day")	AN302
-1d 0hr	90m	PhotoOp/OD/ManPlan (OD and Maneuver Planning for –18hr TCM)	AN303
-23.0hr	210m	High Gain on Earth Telecom Track	
-19.5hr	90m	RCS TCM ("Minus –18hr Hour")	AN304
-18.0hr	90m	PhotoOp/OD/ManPlan (OD and Maneuver Planning for –12hr TCM)	AN303
-17.0hr	210m	High Gain on Earth Telecom Track	
-13.5hr	90m	RCS TCM ("Minus –12hr Hour")	AN305
-12hr	90m	PhotoOp/OD/ManPlan (OD and Maneuver Planning for –6hr TCM)	AN303
-11hr	270m	High Gain on Earth Telecom Track (Last Ground Intervention Opportunity)	
-6.5 hr	90m	RCS TCM ("Minus –6hr Hour")	AN306
-5.0hr	75m	PhotoOp/OD/RSEN Init	AN307
-5.0hr	Continuing	Low Gain Track, S/C on Target	
-3.5hr	90m	RCS TCM ("Minus –3hr Hour")	AN308
-2.0hr	30m	PhotoOp/OD (10m P.O., 20m OD)	AN309
-1hr 30m	90m	Encounter Sequence	SEQ50
-1hr 30m	10m	PhotoOp	Do.
-1hr 15m	10m	PhotoOp	Do.
-55m	25m	OD	Do.
-27m	27m	RSEN	Do.
-5m	2.5m	1st Close Approach Sequence	SEQ51
-2.5m	1.5m	2nd Close Approach Sequence	SEQ52
-90s	65s	3rd Close Approach Sequence	SEQ53
-25s	25s	4th Close Approach Sequence	SEQ54

The rehearsal was very successful overall. All AutoNav operations succeeded:

- Execution of rehearsal RCS TCM #1, a 2400-km B-Plane deflection, or 1.7 m/s, was normal, with performance (determined afterward from radio data) to be within 1.5 percent.
- Onboard simulation of images, PhotoOp operations, including image processing, OD, and maneuver planning for RCS TCM #2 occurred normally.

- Execution of rehearsal RCS TCM #2, a 500-km 0.3-m/s burn, was normal.
- Entry into RSEN mode was normal. RSEN improves position knowledge to better than 0.5 km in the field and 5 seconds downtrack.
- Late image processing allowed RSEN to track Spoof to within 30 seconds of encounter, and the approach late-encounter sequences were initiated within their expected uncertainties.

5.2.2.2 The Encounter. Perhaps the most challenging aspect of the encounter to AutoNav was the lateness of expected acquisition of the target in the images. Had the approach exposures not been limited to 5 seconds or less due to scattered light and light leakage into and within MICAS, Braille would likely have been imaged in time for the E-5-day TCM and possibly the E-10-day (50- to 100-second exposures would have been taken). As it was, the target was not detected until 3 days before closest approach and then only with extensive post-processing on the ground. The AutoNav system only detected a strong enough signal to "lock on" to the asteroid image at –17 hours, again due to the dual limitation of short exposures and scattered light.

As DS1 approached Braille, the normal image-processing operations continued (see Section 5.2.1). These PhotoOps were also used to take the first images of Braille, for use in follow-on OD. Unfortunately, Braille was not seen in the first two attempts (on July 24 and July 25). On July 26, AutoNav again failed to see Braille, but with intensive image processing on the ground, including picture addition, an extremely faint "phantom" appeared, approximately 350 km from the nominal expected position of Braille. This represented about a 2-sigma error from the recently delivered Braille ephemeris. Because of this somewhat large apparent ephemeris change (based on suspect Earth-based data), and the fact that the radio solution was indicating that the E-5-day TCM had performed nominally, it was decided to cancel the E-2-day TCM. In other words, aside from the apparent ephemeris error, which was not nearly well enough determined by the "phantom" observation to act upon, there was no reason to implement the maneuver.

On July 27, three reliable but very dim images of Braille were acquired and processed on the ground. The observed position of Braille was consistent with the earlier "phantom." From these, a design was constructed for the E-1-day TCM. Using the AutoNav software on the ground, as would have been onboard if a higher signal had been available from MICAS, a maneuver file was created that included the TCM. This file was uplinked for use in the execution of the E-1-day TCM (see Figure 5-9). That same day, from 18:30 to 21:00, the E-1-day TCM successfully executed. On July 28, Braille was not yet bright enough for AutoNav to lock on, but ground processing was able to extract another two detections from the downlinked images. These indicated that the spacecraft was sufficiently on target to warrant cancelation of the E-18-hour TCM.

During the next PhotoOp of July 28, 18 pictures of Braille were scheduled and taken. An unknown number of these images locked on. After image processing, AutoNav attempted to store the processed images into the OD file. A previously unknown software fault in AutoNav caused the vector of stored planning cycles to be exceeded by 1. This caused a memory write out-of-bounds and a subsequent reboot. However, three pictures had been scheduled for downlink. From later examination of these images that were subsequently downlinked, it seemed reasonable to assume that most of the other pictures were successfully processed. The recovery of DS1 from safing (in which it found itself after the reboot) was accomplished in little more than three hours.

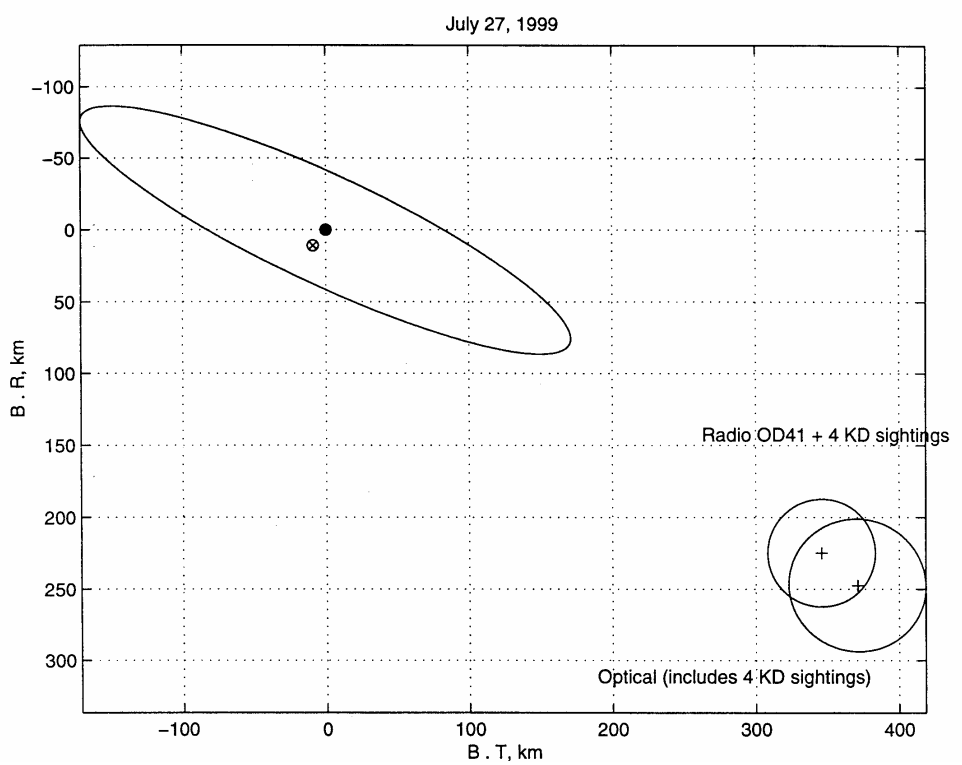

Fig. 5-9. Pre-E-1-day TCM "flight OD" Braille B-plane.

Amidst the recovery efforts, operations continued with the preparation for the E-6-hour TCM. With the three pictures that had been received, the AutoNav team completed the operation interrupted onboard by the software fault, but with much less data. The optical data indicated that the E-1-day TCM had successfully placed the spacecraft within 25 km of Braille, but not on the desired shadowed side. A maneuver was designed to place the spacecraft on a 15-km impact parameter trajectory. However, the solution was chosen from the distribution of solutions such that the target point would be biased to the outside. In other words, with the 1-sigma variance of solutions at 10 km, it was decided that an extra margin of safety was warranted and taken. This Maneuver file was created and uplinked shortly before the spacecraft turned away from Earth for the E-6-hour TCM (See Figure 5-10). The TCM executed normally on July 29, at 22:25 UTC.

AutoNav continued to image Braille and perform OD using the resultant data down to E-30 minutes. At that point, AutoNav switched over to make use of the APS sensor (to avoid making use of the CCD, which was expected to grossly overexpose and bleed poorly to the asteroid during the encounter.) Unfortunately, no signal from Braille was observed above the AutoNav APS threshold. The following is an outline synopsis of the Nav events of the last 30 minutes.

- At closest approach minus 27 (ACA-27) minutes: RSEN activated; AutoNav switched to APS sensor.
- ACA-20 minutes: An unknown signal (probably a cosmic ray) spoofed AutoNav into a one-quarter APS field-of-view (FOV) correction. Braille remained in the APS and CCD fields; however, no frames were preserved.

- Down to ACA-3 minutes: Braille was in APS and CCD fields; no science frames were taken or preserved. Nav activated the first encounter sequence based on *a priori* data. Sequences were scheduled for ACA-300, -150, -90, and -25-second initiations.
- ACA-150 seconds: The first CCD science frame was taken. Braille was barely out of MICAS CCD FOV due to picture editing; outside of all subsequent picture APS and CCD fields.
- Inside 20 seconds: Braille was imaged in the infrared FOV.
- ACA-10 seconds: The sequence stopped tracking Braille pictures inbound.
- ACA+15 minutes: DS1 was back on the nominal (e.g., pre-flyby ephemeris) Braille track. The first successfully taken and returned close-up images of Braille occurred here. APS images showed an extraordinarily dim image—10 DN, with 1000 DN expected. CCD images showed 400 DN, one-tenth "fullwell," with 1/2 to 1 expected and that on the "bright side."
- Post-encounter reconstruction indicated that approach Braille images were 1 to 2 magnitudes dimmer than outbound, perhaps due to the presented geometry of the irregular figure of Braille. Outbound images were also very dim, by factors of 5 to 10 than expected.

From the above timeline, it is apparent that the close-approach events did not proceed according to plan. In review, there was insufficient signal in the APS detector to allow AutoNav to detect Braille. Figure 5-11 shows diagrammatically the expected and received Braille signal on approach. Because no signal from Braille came above the minimum threshold, RSEN never locked on. One of the principal causes of the lack of detection was the poorly characterized nonlinearity of the APS detector. This nonlinearity in the camera response is shown in

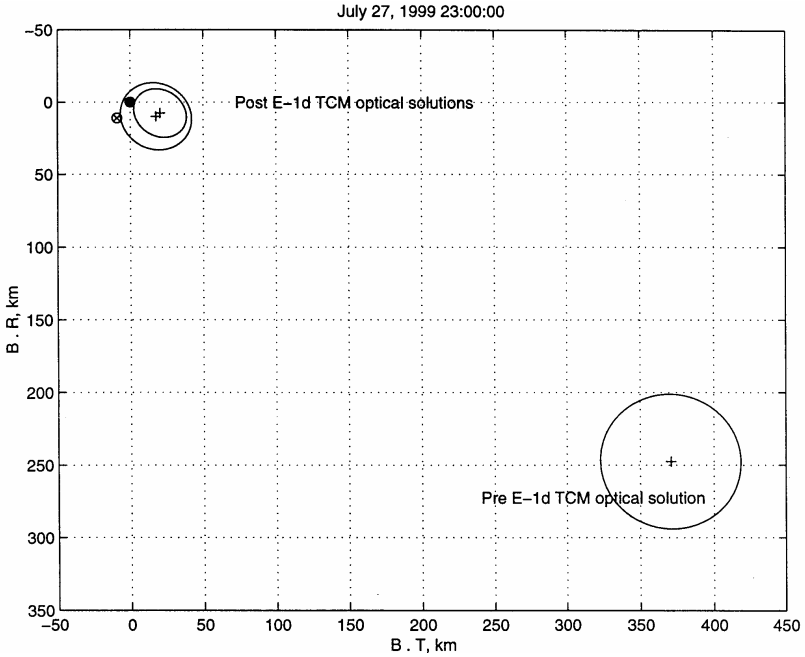

Fig. 5-10. Pre-E-6-hour TCM B-plane, July 27.

Figure 4-6. Additionally, a noise spike, presumed to be a cosmic ray, penetrated the threshold, and AutoNav temporarily locked onto this, causing a deflection in the trajectory. Figure 5-12 shows the effect of this deflection on the position of Braille in the two visual fields of view versus the nominal trajectory that would have been followed if there had not been the cosmic ray event.

5.2.2.3 Results of Encounter. Despite the fact that the performance of the system during the Braille flyby was thwarted, it is nevertheless the case that operability and accuracy of the AutoNav close-approach system had been demonstrated in the testbeds, and more importantly, in-flight during the rehearsal. This was demonstrated using the few acquired images of Braille post-encounter. When these were provided to RSEN, accurate solutions of the spacecraft position were obtained with just 1 CCD image, leading to the unavoidable conclusion that had this detector been used, the encounter would likely have been very successful. Figure 5-13 shows the B-plane results of this analysis.

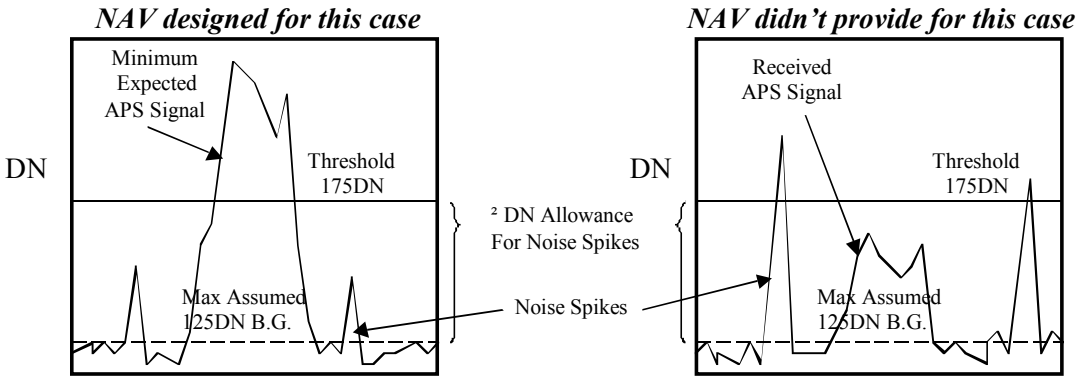

Fig. 5-11. Diagrammatic view of received RSEN signal.

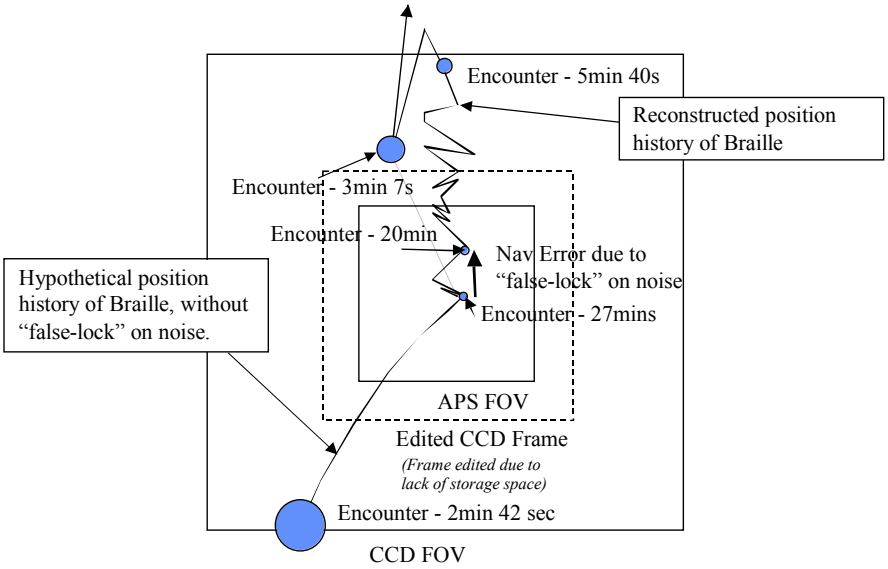

Fig. 5-12. Reconstructed nominal vs. perturbed Braille field-of-view flight path.

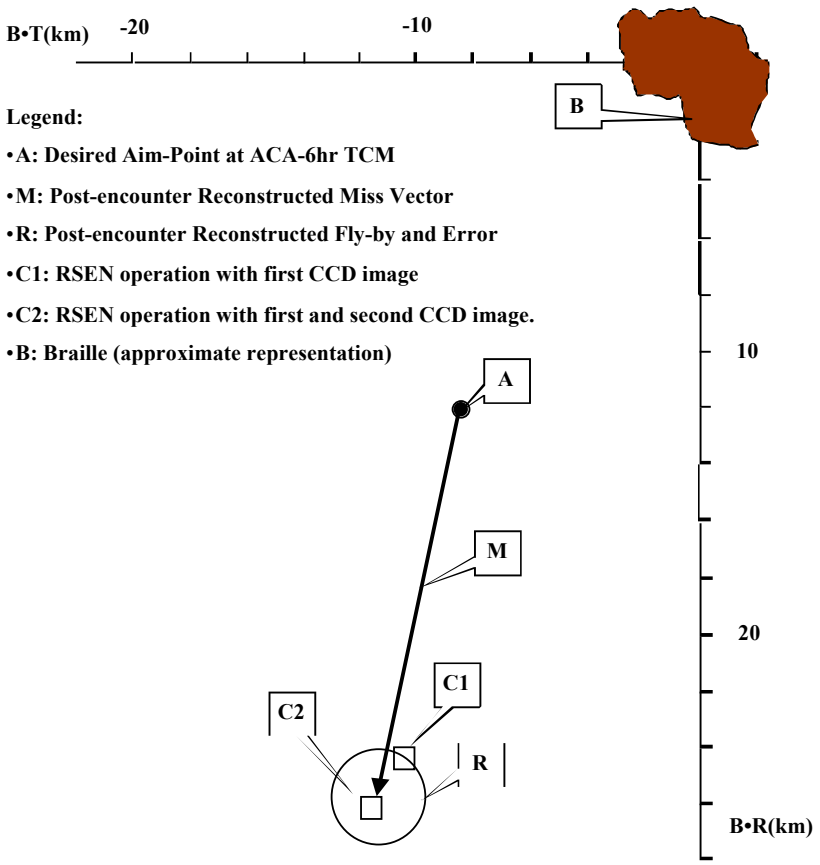

Fig. 5-13. Encounter results using post-encounter CCD Braille pictures.

5.2.3 Post-Braille Cruise Operations

Within a few weeks after the Braille encounter, navigation events began again in earnest. In order to achieve the targeting requirements for an encouner with asteroid Wilson–Harrington in January of 2001, it was necessary to start operating the main (IPS) engine within days of closest approach. Fortunately the desired thrust attitude was close to the attitude of the spacecraft with its high-gain antenna oriented on Earth, and so it was possible to take advantage of the extensive scheduled DSN tracking scheduled after the Braille encounter while thrusting. The first post-Braille PhotoOp navigation event took place on August 9. HGA-on-Earth operation of the IPS continued, with another PhotoOp on August 16 and August 23. The first two of these PhotoOps were very successful. However, the third revealed another coding flaw in AutoNav, where due to a dearth of sufficiently bright targets and the need to "double-up" on a single good target at an imaging opportunity, an internal array was overrun, causing the spacecraft to safe. With the real (as opposed to opportunistic HGA-on-Earth) IPS thrusting scheduled to start on that day, a rapid spacecraft recovery took place, and the mission burn began early on August 25. With the Navigation Team focused on accomplishing the next eight weeks of thrusting and assuring the safety of OpNav events, a one-month hiatus in PhotoOps was declared. Starting on September 20, PhotoOp events began again, and for seven weeks, these were weekly events. There was also a change of strategy. It was decided for a number of reasons (not the least of which was sheer exhaustion of the team) that to simplify AutoNav operations, picture planning would revert to the original design. That is, optical frames would be boresighted on the asteroid target (actually the targets had to be substantially offset from the center of the field in the CCD

due to large, severely attenuating scratches in the optics at that point), and the system would acquire any stars available. This substantially reduced the "manhandling" of the system and allowed it to operate largely autonomously.

Figure 5-14 shows the post-fit residuals for one of the final AutoNav solutions with a data arc extending from September 27 to November 1. These make an interesting comparison to Figure 5-3, showing a factor of 2–3 improvement in image-processing performance with a drastic reduction in effort. In fact, the effort was literally reduced to zero. For the period shown in Figure 5-14, the spacecraft was navigating itself, with no updates or changes to its process. This turned out to have substantial advantages; with several critical programs operating (and experiencing navigational problems), the DS1 Radio Navigation Team was released to concentrate on these challenges, while DS1 navigated itself, which is perhaps the best characterization of the validation of AutoNav.

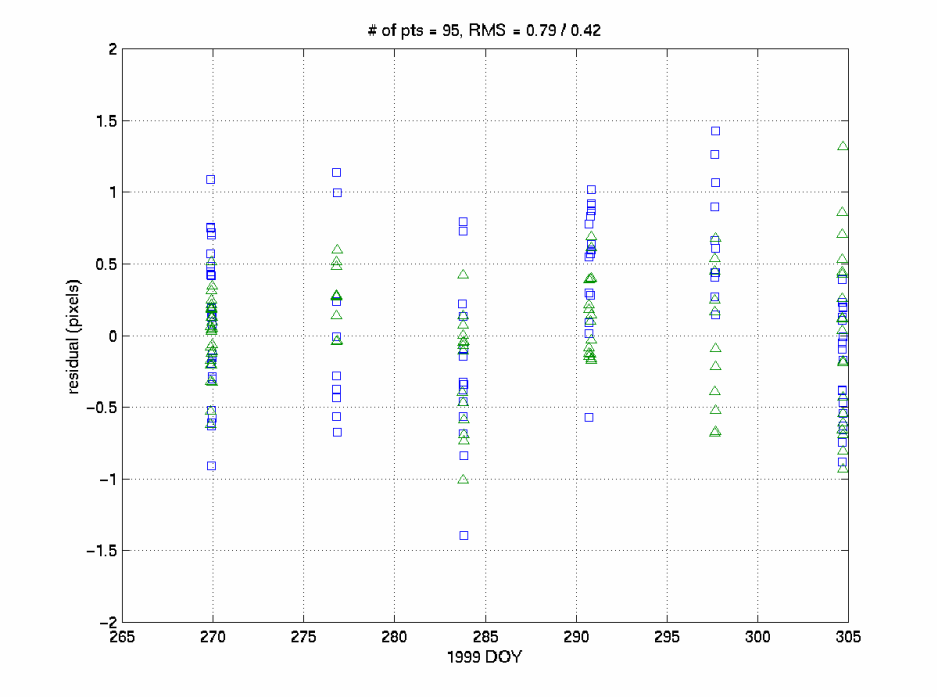

Fig. 5-14. Post-Braille AutoNav data arc and residuals.

Section 6

Observations, Lessons Learned, and Future Applications

The following section describes lessons learned by the Navigation team over the course of the DS1 primary mission. These include the radio navigation and optical observations, as well as the *a priori* navigation requirements and their effects on the applicability of AutoNav to future missions.

Succinct descriptions of lessons learned on a wider (project) level can be found in [10].

6.1 Radio Navigation

6.1.1 Modeling and Predicting Nongravitational Forces

6.1.1.1 Unbalanced Reaction Control System. With DS1, the use of an unbalanced Reaction Control System for attitude control during the coast phases of the mission made DS1 a difficult spacecraft to navigate. Painstaking modeling of spacecraft accelerations allowed for adequate results to be obtained. The use of a balanced RCS or momentum wheels for future spacecraft development is strongly encouraged.

6.1.1.2 Uncertainties in the Performance of the Ion Propulsion System. When navigating a low-thrust mission, the navigation plan should account for large uncertainties in the predictability of the expected thrust. It should also account for possible degradations in the thruster performance and/or the power source (typically, the solar arrays).

Future missions that plan a short coast (i.e., non-thrusting) arc before an encounter need to be studied carefully to make sure the thrust predictability and orbit determination requirements are consistent with a short encounter timeline.

6.1.2 Tracking Strategy

The large uncertainties in the accelerations that accompanied DS1 highlighted the navigational insensitivity of a single pass of Doppler and range tracking data to geocentric declination. For spacecraft that are constrained to using short, infrequent tracking passes, this can be ameliorated by the use of consecutive tracking passes, consecutive ranging passes from ground stations in alternate hemispheres, or very long baseline interferometry (VLBI) measurements.

6.2 AutoNav Lessons Learned—The Braille Encounter

6.2.1 Imaging Instrumentation

The problems with the MICAS instrument that were described earlier point out the need for greater understanding and more methodical analysis of the expected response of the imaging instrument(s) to the target. In flight, this will amount to a need for more thorough pre-flight

analysis of the response curves in the imaging instruments and a very thorough imaging campaign of the flyby target in the days leading up to an encounter.

6.2.2 Algorithm Reviews

As a result of the software error noted in Section 5.2.2.2 and a similar fault in August 1999, an extensive review of the AutoNav code was undertaken by non-Navigation team members, which fortunately revealed only two or three additional problems, none nearly so dramatically serious. This underscores the need to encourage "non-peer" peer groups to perform pre-flight code examinations and test plan reviews.

6.2.3 File Storage Space and Memory Allocation

In the DS1 file system, there was extremely limited space onboard for stored images. This points out a need for a more carefully thought out science return strategy that might be undertaken with less optimistic expectations. Reallocation of random access memory (RAM) space might have been possible but was not undertaken, nor was taking and preserving earlier, more reliable, but less resolved images. This might have returned more science data at the expense of engineering and housekeeping telemetry.

6.3 Future AutoNav Adaptations and Developments

6.3.1 Algorithm and Software

The DS1 AutoNav demonstration was a very low-cost initial foray into the uncharted technological territory of autonomous deep-space navigation. As such, it was very successful, but there is much work left to be done. The Deep Impact mission, scheduled for launch in January 2005, has taken AutoNav one step further by making the system robust enough for the mission to place 100 percent reliance upon it for the principal mission science objectives. Further developments are necessary before AutoNav can be reliably placed onboard a future mission and flown with minimal attention. Automated picture planning will need to be added, so that AutoNav can determine its own imaging data needs. Extended body modelling and image processing will need to be added to allow for Earth–Moon departure navigation, large-body approach Nav, and orbit insertion navigation. Landmark modelling and tracking should be added for orbital operations. Increased flexibility and intelligence with regard to maneuver planning would aid orbital operations as well. All of these improvements will allow future missions such as Jupiter Icy Moons Orbiter (JIMO) to confidently use AutoNav.

6.3.2 AutoNav Hardware

Along with planned improvements in the AutoNav software, developments in navigation-related hardware will soon be linked to AutoNav. The 2005 Mars Reconnaisance Orbiter (MRO) mission will fly a specifically designed Navigation camera. With a high-sensitivity CCD detector, 7-cm aperture, and 1.4-deg FOV, this camera offers ideal characteristics to perform optical navigation in many different circumstances. Adding a gimbal to the camera to allow navigation imaging without disrupting normal spacecraft activities will be a powerful enabler of accurate low-cost navigation, as will providing a dedicated central processing unit (CPU) for navigation purposes. Having a Nav computer will greatly reduce the intricate software interactions that challenged the DS1 and Deep Impacet (DI) engineers, and eliminate the very

difficult issues of computational throughput and data transfers from the cameras. All of these hardware improvements, as well as the software issues mentioned above, are planned to be addressed aboard the Mars Telesat Orbiter (MTO) mission, scheduled for launch in 2009. This mission will fly a dedicated AutoNav hardware package including a Nav CPU and gimballed MRO-style cameras. Additionally, a wide-angle camera will be added, and AutoNav will have its capabilities further increased to be the means of effecting a rendezvous. MTO will carry a small (20-cm) sphere that is a model of the proposed Mars Sample Return (MSR) Orbiting Sample (OS), which will be lofted from the surface of Mars carrying 500 g of Mars rocks and soil. This OS will be detected and tracked by AutoNav. Treating the OS like a small moon, AutoNav will allow MTO—and eventually MSR—to maneuver close to the OS, and in the case of MSR, to capture the sample and return it to Earth. The navigation experiment on DS1 was difficult and had moments that seemed less than successful. Nevertheless, it achieved very good overall success, and that success is definitely feeding forward to future missions.

References

[1] Rayman, M. D. Varghese, P., Lehman, D. H., Livesay, L. L., "Results from the Deep Space 1 Technology Validation Mission," IAA-99-IAA.11.2.01, *50th International Astronautical Congress*, Amsterdam, The Netherlands, Octgober 4-8, 1999.

[2] McElrath, T. P., Lewis, G. D., "Ulysses Orbit Determination at High Declinations," *Proceedings of the Flight Mechanics/Estimation Theory Symposium 1995*, Greenbelt, Maryland, May 1995, pp. 277-287.

[3] Hamilton, T. W., Melbourne, W. G., "Information Content of a Single Pass of Doppler Data from a Distant Spacecraft," *JPL Space Programs Summary 37-39*, Vol. 3, May 1966

[4] Wolff, P. J., Pinto, F., Williams, B. G., Vaughan, R. M., "Navigation Considerations for Low-thrust Planetary Missions," *AAS/AIAA Space Flight Mechanics Meeting*, Monterrey, California, February 1998.

[5] Vaughan, R. M., Riedel, J. E., Davis, R. P., Owen, W. M., Synnott, S. P., "Optical Navigation for the Galileo Gaspra Encounter," AIAA paper 92-4522, *AIAA/AAS Astrodynamics Conference*, Hilton Head, South Carolina, August 1992.

[6] Bhaskaran, S., Desai, S. D., Dumont, P. J., Kennedy, B. M., Null, G. W., Owen, W. M., Riedel, J. E., Synnott, S. P., Werner, R. A., "Orbit Determination Perforamnce Evaluation of the Deep Space 1 Autonomous Navigation System," *AAS/AIAA Space Flight Mechanics Meeting*, Monterrey, California, February 1998.

[7] Coverston-Caroll, V., Williams, S. N., "Optimal Low Thrust Trajectories Using Differential Inclusion Concepts," *AAS/AIAA Spaceflight Mechanics Meeting*, Cocoa Beach Florida, February 1994.

[8] Riedel, J. E., Bhaskaran, S., Synnott, S. P., Desai, S. D., Bollman, W. E., Dumont, P. J., Halsell, C. A., Han, D., Kennedy, B. M., Null, G. W., Owen Jr., W. M., Werner, R. A., and Williams, B. G. "Navigation for the New Millenium: Autonomous Navigation for Deep Space-1", *Proceedings of the 12th International Symposium on Flight Dynamics*, Darmstadt, Germany, June 1997.

[9] Campbell, J. K., Jacobson, R. A., Riedel, J. E., Synnott, S. P., Taylor, A. H., "Voyager I and II Saturn Encounter Orbit Determination," AIA-82-0419, AIAA 20th Aerospace Sciences Meeting, January 11–14, 1992, Orlando, Florida.

[10] Taylor, J., Fernández, M. M., Bolea Alamañac, A. I., Cheung, K.-M., *Deep Space 1 Telecommunications*, DESCANSO Design and Performance Summary Series, Article 2, Jet Propulsion Laboratory, Pasadena, California, October 2001.

Abbreviations and Acronyms

ACA	at closest approach
ACS	Attitude Control Subsystem
APE	Attitude Planning Expert
APS	active pixel sensor
AutoNav	Autonomous Onboard Optical Navigation System
CCD	charge-coupled device
CPU	central processing unit
DDOR	delta differential one-way range
delta-V	change in velocity
DESCANSO	Deep Space Communications and Navigation Systems Center of Excellence
DI	Deep Impact
DS1	Deep Space 1
DSN	Deep Space Network
E	encounter
EEPROM	electrically erasable programmable read-only memory
FK	FrankenKenny
FOV	field of view
FP	fault protection
FSW	flight software
HGA	high-gain antenna
IAT1	IPS Acceptance Test 1
IAT2	IPS Acceptance Test 2
IPS	Ion Propulsion System
JIMO	Jupiter Icy Moons Orbiter
JPL	Jet Propulsion Laboratory
Ka-band	deep-space frequency band: 31.9 to 32.1 GHz (down)
LGA	low-gain antenna
LOS	line of sight
ManPlan	Maneuver Planning
MB	megabyte
MICAS	Miniature Integrated Camera and Spectrometer
MRO	Mars Reconnaisance Orbiter
MSR	Mars Sample Return
MTO	Mars Telesat Orbiter
NASA	National Aeronautics and Space Administration
Nav	navigation
NavRT	Nav Real-Time
NMP	New Millennium Program
nongrav	nongravitational
OD	orbit determination
opnav	optical navigation
OS	Orbiting Sample

RAM	random access memory
RCS	Reaction Control System
RSEN	Reduced State Encounter Navigation
SRP	solar radiation pressure
SRU	Stellar Reference Unit
Starcat	AutoNav's Star Catalog file
TCM	trajectory correction maneuver
TVC	thrust vector control
VLBI	very long baseline interferometry